

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of: Sunil Podar et al.
Serial No.: 09/745,909
Filing Date: December 21, 2000
Examiner: Aravind K. Moorthy
Group Art Unit: 2131
Title: **METHOD AND SYSTEM FOR
AUTHENTICATED ACCESS TO INTERNET
PROTOCOL (IP) MULTICAST TRAFFIC**

MAIL STOP AMENDMENT
Commissioner for Patents
PO Box 1450
Alexandria, VA 22313-1450

Dear Sir:

DECLARATION PURSUANT TO 37 C.F.R. § 1.131

We, the undersigned, hereby declare and state that:

1. We are each over the age of 21 years, of sound mind, and competent in all respects to make this Declaration.
2. We are the Inventors of the subject matter of the above-referenced application, entitled *Method and System for Authenticated Access to Internet Protocol (IP) Multicast Traffic*, filed December 21, 2000 (the "Application").

3. The Examiner rejected Claims 1, 2, 4-17, 19-32, and 34-47 of the Application in an Office Action mailed August 17, 2006, based, in whole or in part, on U.S. Patent No. 6,963,573 to Cain et al. ("*Cain*") filed on September 13, 2000 (the "Effective Date") and issued on November 8, 2005.

4. We described certain aspects of the subject matter of the Application in disclosure documents and implemented certain aspects of the subject matter of the Application in an operating computer program (the "Program"). Exhibit A includes a first disclosure document disclosing certain aspects of the subject matter of the Application, which existed prior to the Effective Date. Exhibit B includes a second disclosure document disclosing certain aspects of the subject matter of the Application, which existed prior to the Effective Date. Exhibit C includes an agenda for an executive briefing that included the Program's demonstration, which occurred prior to the Effective Date. The submission and modification dates of the disclosure documents and the agenda have been redacted. The Program operated on a computer at the offices of Cisco Systems, Inc. in San Jose, California.

5. The Program operated prior to the Effective Date for authenticated access to multicast traffic by: receiving an Internet group management protocol request at an access network router operable to authenticate a plurality of requests received from a plurality of customer premise systems, the received request identifying a user requesting to join an IP multicast channel, the IP multicast channel selected from a bundle of IP multicast channels offered for receipt by the user as a multicast package on a subscription basis; authenticating access privileges of the user to the multicast channel; and disallowing the request in response to at least an unsuccessful authentication.

6. The Program operated prior to the Effective Date using a system for authenticated access to multicast traffic that included at least: means for receiving an Internet group management protocol request at an access network router operable to authenticate a plurality of requests received from a plurality of customer premise systems, the received request identifying a user requesting to join an IP multicast channel, the IP multicast channel selected from a bundle of IP multicast channels offered for receipt by the user as a multicast package on a subscription basis; means for authenticating access privileges of the user to the

multicast channel; and means for disallowing the request in response to at least an unsuccessful authentication.

7. The Program operated prior to the Effective Date using a system for authenticated access to multicast traffic that included at least: logic encoded in media; and the logic operable to receive and authenticate a plurality of requests received from a plurality of customer premise systems, at least one of the plurality of requests comprising an Internet group management protocol request for a user to join an IP multicast channel selected from a bundle of IP multicast channels offered for receipt by the user as a multicast package on a subscription basis, to authenticate access privileges of the user to the multicast channel and to disallow the request in response to at least an unsuccessful authentication.

8. The Program operated prior to the Effective Date for providing premium content services over a network using Internet protocol (IP) multicast channels by: provisioning user access privileges to an IP multicast channel providing premium content, the premium content including at least one of video, audio and data; authenticating access privileges of a user to the IP multicast channel upon receiving an Internet group management protocol request at an access network router operable to authenticate a plurality of requests received from a plurality of customer premise systems, the received request identifying a user requesting to join an IP multicast channel to receive the premium video content, the IP multicast channel selected from a bundle of IP multicast channels offered for receipt by the user as a multicast package on a subscription basis; and disallowing the request in response to unsuccessful authentication.

9. The Program operated prior to the Effective Date for providing premium content services over a network using Internet protocol (IP) multicast channels by: receiving an Internet group management protocol request at an access network router operable to authenticate a plurality of requests received from a plurality of customer premise systems, the received request identifying a user requesting to join an Internet protocol (IP) multicast channel; authenticating access privileges of the user to the IP multicast channel by at least one of: determining whether the IP multicast channel is a public multicast channel; determining whether the user is logged in to a service provider providing a service including the IP multicast channel; determining whether the user is logged in to the service including the IP multicast channel; successfully authenticating access privileges of the user to the IP multicast channel in response to at least one of determining the multicast channel is a public multicast channel and determining the user is logged in to the service provider and the service; unsuccessfully authenticating access privileges of the user to the IP multicast channel in response to at least one of determining the user is not logged in to the service provider and determining the user is not logged in to the service; terminating the request in response to at least an unsuccessful authentication; and processing the request in response to at least a successful authentication.

10. I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true. Further, I declare that these statements are made with the knowledge that willful false statements, and the like so made, are punishable by fine or imprisonment, or both, under Section 1001, Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the Application or any patent issuing thereon.


ATTORNEY DOCKET NO.
062891.0505

PATENT APPLICATION
09/745,909

5

Declaration pursuant to 37 C.F.R. § 1.131 in regard to 09/745,909.

Signed this 13th day of November, 2006.



Sunil (nmi) Podar

Signed this 7th day of NOVEMBER, 2006.



Sunil K. Chandrupatla

Signed this ~~13th~~ day of ~~November~~, 2006.

Sandeep (nmi) Salsena

Signed this _____ day of _____, 2006.

Kali Prasanna Mishra

Signed this _____ day of _____, 2006.

Sampath Kumar Sthothra Bhasham

ATTORNEY DOCKET NO.
062891.0505

PATENT APPLICATION
09/745,909

5

Declaration pursuant to 37 C.F.R. § 1.131 in regard to 09/745,909.

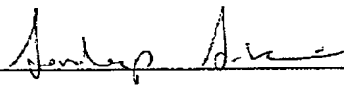
Signed this _____ day of _____, 2006.

Sunil (nmi) Podar

Signed this _____ day of _____, 2006.

Sunil K. Chandrupatla

Signed this 07 day of Nov, 2006.



Sandeep (nmi) Saksena

Signed this _____ day of _____, 2006.

Kali Prasanna Mishra

Signed this _____ day of _____, 2006.

Sampath Kumar Sthothra Bhasham

DAL01:929337.1

ATTORNEY DOCKET NO.
062891.0505

PATENT APPLICATION
09/745,909

5

Declaration pursuant to 37 C.F.R. § 1.131 in regard to 09/745,909.

Signed this _____ day of _____, 2006.

Sunil (nmi) Podar

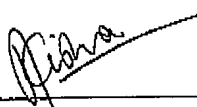
Signed this _____ day of _____, 2006.

Sunil K. Chandrupatla

Signed this _____ day of _____, 2006.

Sandeep (nmi) Saksena

Signed this 15th day of November, 2006.



Kali Prasanna Mishra

Signed this _____ day of _____, 2006.

Sampath Kumar Sthothra Bhasham

ATTORNEY DOCKET NO.
062891.0505

PATENT APPLICATION
09/745,909

5

Declaration pursuant to 37 C.F.R. § 1.131 in regard to 09/745,909.

Signed this _____ day of _____, 2006.

Sunil (nmi) Podar

Signed this _____ day of _____, 2006.

Sunil K. Chandrupatla

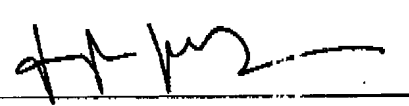
Signed this _____ day of _____, 2006.

Sandeep (nmi) Saxena

Signed this _____ day of _____, 2006.

Kali Prasanna Mishra

Signed this 9th day of November, 2006.



Sampath Kumar Sthothra Bhashara

EXHIBIT A



Document Number: ENG-41119
Revision: G
Author:

Project Manager:

Engineering Information Notice

Tuzigoot

Multicast Video over DSL Architecture

The document describes the network architecture required to provide low to medium bandwidth multicast video service over xDSL.

Contributors

Modification History

Rev	Date	Originator	Comment
A			
B			
C			
D			
E			
F			
G			

Table of Contents

1	Introduction	5
2	Video Multicast Service Description.....	6
3	Document Scope	9
4	Network Architecture For Video Service	10
5	Content Acquisition Architecture.....	11
5.1	Content Format.....	12
5.2	Content Feed	12
5.3	Content Aggregation.....	13
6	Transport Architecture	13
6.1	Transport Network Components.....	14
6.1.1	Transport Core Network Redundancy	14
6.1.2	Multicast Routing Recommendation.....	15
7	Quality of Service (QoS)	16
7.1	Content Aggregation: Quality of Service (QoS)	16
7.2	Transport: Quality of Service (QoS).....	17
7.3	Access: Quality of Service (QoS)	17
8	Access Architecture	18
8.1	General Multicast System Flows.....	21
8.1.1	User/Service Authentication	22
8.1.2	PC issues a IGMP Join Request	22
8.1.3	PC issues a IGMP Leave Request.....	23
8.1.4	User issues a IGMP leave for a channel not authorized to him/her	23
8.1.5	User issues multiple IGMP joins (for multiple channels) without leaving channels	23
8.1.6	Hacker Sends garbage IP multicast streams to a multicast group address.....	23
8.2	Access Architecture with Integrated Routing & Bridging (IRB)	25
8.3	Access Architecture with Route Bridged Encapsulation	26
8.3.1	Introduction.....	26
8.3.2	Protocol Stack	27
8.3.3	Address assignment options for the PCs and CPE.....	27
8.3.4	Configuration Guidelines (6400 NRP)	29
8.3.5	Pros & Cons of Route Bridge Encapsulation mode for providing IP multicast based services	29
8.4	Access Architecture with PPP Over ATM	31
8.4.1	Introduction.....	31
8.4.2	PPPoA Protocol Stack in Access Network	32
8.4.3	Address assignment options for the PCs and CPE.....	32
8.4.4	Configuration Guidelines(PPPoA)	32
8.4.5	Pros & Cons of PPPoA for IP multicast based Services	33
8.5	Access Architecture with PPP over Ethernet.....	35
8.5.1	Introduction.....	35
8.5.2	Protocol Stack	35
8.5.3	Address assignment options for the PCs and CPE.....	36
8.5.4	Configuration for 6400/6100/CPE/PC.....	36
8.5.5	Pros & Cons of PPPoE for providing IP multicast based services	36
8.6	Recommendation (RBE, PPPoA, PPPoE).....	36
8.7	Scalability/sizing: (for RBE, PPPoA, PPPoE cases).....	37
8.7.1	Calculating the total available downstream capacity of 6400:	37
8.7.2	Number of Subscribers that can be supported.....	37
9	SSG based Authentication for IP Multicast group.....	39
10	Technical Issues.....	41
10.1	Channel switching latency and IOS modification/configuration	41
10.2	Enforcing Max Channel Limit.....	41
10.3	IP QoS to ATM CoS.....	42

10.4	Fair DSLAM subtending.....	42
10.5	Smart DSLAM.....	42
10.6	Support WFQ for PPP virtual Access Interfaces	42
10.7	Multicast Enabled DSLAM (Future)	43
11	Appendix A: VLAN Based Architecture for IP Multicast Service	44
12	Appendix B: Filtering Multicast Traffic by ACLs	46
13	Appendix C: VPI/VCI Based CPE Authentication	47
14	Appendix D: IPCP subnet feature:	49
15	Appendix E: Encryption Based Authentication for IP Multicast group	50
16	Appendix F: System Flows.....	50
16.1	System Flows in Route Bridge Encapsulation Architecture	50
16.1.1	CPE boots up/powered down:	50
16.1.2	PC boots up/shuts down/powered down	50
16.1.3	Hacker Tries ARP spoofing	51
16.1.4	User Configures the PC with arbitrary IP addresses	51
16.1.5	Subscriber to Subscriber Communication	51
16.2	System Flows in PPPoA	52
16.2.1	CPE boots up/Powered down.....	52
16.2.2	PC Boots Up/Powered Down.....	53
16.2.3	Hacker Tries ARP spoofing	53
16.2.4	User Configures the PC with arbitrary IP addresses	53
16.3	System Flows in PPPoE	53
16.3.1	CPE boots up/Powered down.....	53
16.3.2	PC Boots Up/Powered Down.....	54
16.3.3	PPPoE Session Establishment.....	54
16.3.4	Subscriber to Subscriber Communication	55
16.3.5	Hacker Tries ARP spoofing	55
16.3.6	User Configures the PC with arbitrary IP addresses	55
17	Appendix G: NRP Configuration Matrix for PPPoE	56
18	Appendix H: NRP Configuration PPPoA,	57
19	Appendix I: NRP Configuration Matrix for RBE	59
20	Appendix J: Changing Existing Architectures for Multicast Video Service	60
21	End of Document.....	62

1 Introduction

Although xDSL¹ networks have proliferated recently, service providers are looking for delivering services with better revenue potential. Such services, while being attractive to customers, must provide sufficient revenue opportunities to be viable. Video multicast service is considered to be a good candidate, and is on the roadmaps of several service providers. Since Internet is not multicast enabled, such services are likely to be provided in service provider's walled gardens. A lot of interest has been generated in having the capability to provide subscription based control of user access to video channels, since it offers an additional revenue model to the service providers, similar to the Cable model.

Tuzigoot project was initiated to prepare and validate an end to end service architecture to provide video multicast service with support for authenticated channel access. The project includes:

- documentation of the network and service architecture,
- laboratory validation of the architecture,
- publishing a white paper on the architecture, and
- modification of IOS (6400 NRP-SSG) in a private branch, to demonstrate SSG based multicast authentication (to be subsequently put on regular IOS Roadmap)

As one of the Tuzigoot deliverables, this document presents xDSL based network architecture (Layer 3 and below) for providing video multicast service, with Service Selection Gateway based access control to video streams. The service level architecture that includes architecture above layer 3 has been separately described in the following document:

Tuzigoot System Functional Specifications (ENG-41377)

http://wwwwin-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/FS_SystemFsTuzigoot.doc

The next section describes the video multicast service in general, and some available third party services/products that can be used for providing video services. Subsequent chapters present the detailed network architecture and includes the content acquisition, transport, and different applicable flavors of DSL access architectures (Route Bridge Encapsulation, PPPoA, and PPPoE) with their pros and cons. The architecture focuses on video service delivery, although it can be used to deliver other services (telecommuting, Internet) as well.

¹ Although this document may have specific references to ADSL, the described architecture is also applicable to other flavors of xDSL (IDSL, SDSL, VDSL).

2 Video Multicast Service Description

Although video is being delivered to DSL users in a variety of ways today, they have met with limited acceptance, due to various reasons such as poor picture quality, low frame rate, and small picture size/excessive bandwidth requirements. These deficiencies are further aggravated by the network's limitation in providing reliable transport and multicast deployment. On the other side, those that do receive a high quality video signal need DSL connections that can deliver upwards of 5-6Mbps of bandwidth. This excludes the majority of DSL users, since only a small percentage of DSL subscribers today have connections with bandwidth higher than 1 Mbps. Therefore there is a need for a solution that will enable delivering quality video service to the end user(s) with lesser bandwidth.

Some content providers (e.g., CoolCast) are already addressing this market by providing premium TV channels at lower bandwidth ranges (e.g., 300Kbps, 600Kbps). These channels are commercially available to service providers for multicasting. Additionally, technology is now available (Clear Band, for example) to multicast quality video streams at lower bandwidth. Defining an end to end architecture for delivering such video services, therefore, has assumed greater importance at present.

Tuzigoot project has been initiated to define such an end-to-end architecture based on currently available xDSL technology and other third party solutions (such as Coolcast and Clearband), and to recommend any additional development required to provide a viable video delivery system that can be successfully deployed in a Service Provider's xDSL network.

Some sample third party solutions for video multicast service are described below (other solutions may exist).

CoolCast Service and Architecture

CoolCast is a video content provider that delivers video channels (broadcast/cable channels such as CNN, Disney, Bloomberg) in a format suitable for providing video to PC Internet users. CoolCast aggregates video content from several sources and offers them to service providers as IP multicast streams via its private satellite network. The Service Providers receive the multicast video streams from CoolCast (on commercial basis), and multicast the content in their own networks. CoolCast simplifies content acquisition for a Service Provider interested in deploying video multicast service.

Apart from video channels traditional content producers, CoolCast ability to deliver multiple channels to the PC enables it to provide a broad variety of contents, ranging from traditional basic service to video targeted towards special interest groups, pay-per-view, distance learning, or product-specific marketing.

CoolCast currently provides the multicast video channels at rates (300 and 600Kbps) that cover most xDSL subscribers today. These rates enable users to watch high frame rate video in quarter or half-screen size windows². A one-time free download of the CoolCast "Plug-In" software and a standard media player enables the user to view CoolCast video. Beyond this no new consumer action is required. Users request and view web pages normally, using standard web browsers

² CoolCast has facility to simulcast a video channel at 8 different bit rates, to provide video quality commensurate to a user's xDSL-link bandwidth.

CoolCast delivers video streams over its private satellite based CoolCast Network. This network augments the existing Internet by enabling video to travel directly to Service Providers via high-speed satellite communications, separate from normal Internet data. The CoolCast "overlay" network bypasses the normal Internet and provides dedicated capacity for each video stream. This enables CoolCast video to avoid the unpredictability of the current Internet's inherent "party line" nature. Text and graphics associated with the video (information related to the video, or simply advertisements) are received by the PC from the Internet and displayed to the user along with the video. The consumer sees video and the associated text/graphics inside a web page, without being aware of the separate video and web page transport paths.

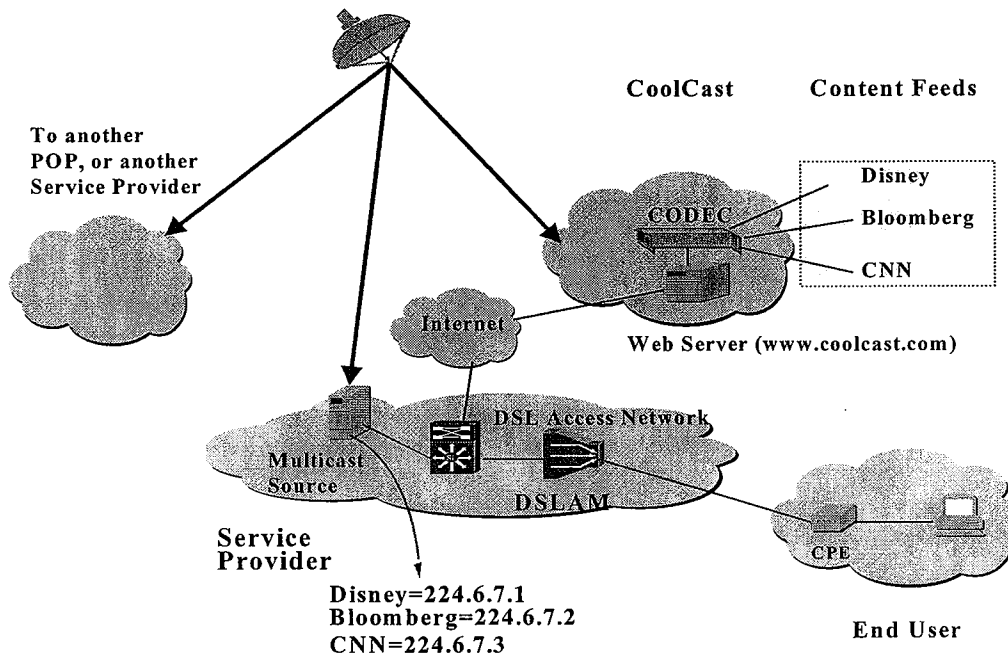


Figure 1 The CoolCast Service Model

ClearBand

ClearBand has created a video server product that greatly reduces the bandwidth requirements for delivering high quality video to the end user. The video source (e.g., video camera, VCR) is connected via RCA jacks to an Intel video interface card on a PC running the Clearband software. The software encodes the input video stream into an MPEG2 compatible format at user specified bit rates (300 kbps onwards) and multicasts the data as an IP multicast stream. A thin client application (Microsoft 95/98/NT supported) is required to decode and view the video stream. However, the plug-in need not be downloaded nor stored on the PC by the user; the ClearBand Web server delivers the plug-in along with the IP multicast stream, resulting in a more user-friendly multicast service.

The major advantage, however, is that the proprietary ClearBand compression algorithm generates higher quality video requiring significantly less bandwidth (500kbps provides reasonably good quality full screen video).

Tuzigoot Project:

Tuzigoot project is to be conducted in two phases. Phase 1 includes the specification of the end to end reference architecture for video service delivery, modification of NRP-SSG on a private branch to support authenticated access to multicast groups, and lab demonstration of the system. This will support low to medium bandwidth content feeds (from Coolcast, Clearband generated content, Webcam) of less than 1Mbps bandwidth. Phase 2 is being envisaged to be an integrated solution involving third party vendors as required and will address issues discovered in phase 1. More details can be found in the Tuzigoot PRD at

http://wwwwin-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/PRD_Tuzigoot.doc

Although the primary focus of Tuzigoot architecture is to deliver low to medium bandwidth IP multicast video streams, it can be used for deploying other video-based services such as:

- Webcam based applications
- Home security applications
- Community of interest video multicasting
- Distance learning (within the service provider walled garden)

The user is expected to use his/her PC to view the service; a set top box (and TV) are not required.

Some associated requirements for Tuzigoot architecture are described below:

- Deliver Multicast content (e.g., Coolcast) to residential users with DSL connections of 384Kbps, 640Kbps, and 1.5Mbps (G.lite maximum). The multicast channel feeds from Coolcast will be at 300Kbps and 600Kbps rates. Feeds from Clearband can be at 300K, 500K, 1Mbps, and 1.5 Mbps.
- Support a mixed access bandwidth environment where, for example, some users will watch 300Kbps video streams while other may watch 600Kbps streams
- Provide enough capacity for up to 100 channels (allowing for other local content).
- Though the initial focus is on Coolcast, the architecture MUST consider additional feed sources such as locally generated content, Webcam, etc.
- The architecture should address:
 - Scalability (greater number of channels, growing service deployment, etc.)
 - Multi-PC access from home
 - Security (network address hijacking issues, ARP replies, etc.)
- Document the architecture, pros, cons, and limitations for each feasible architecture
- Develop, evaluate and recommend architectures for access network; evaluate 3 options: PPPoA, Half Bridging, or PPPoE.

The project entailed proof-of-concept development of Multicast Authentication on a private branch of the 6400/Vulcan image³ (on successful demonstration, the same feature may be made available in IOS later on).

³ This has been completed as of 11/19/99. These changes to IOS 12.05(DC) on a private branch, is available for demonstration.

3 Document Scope

This document describes the network architecture for the Tuzigoot project at layer 3 and below. The service architecture (above layer 3) to provide the end to end video service is described in the following document:

Tuzigoot System Functional Specifications (ENG-41377)

<http://www.in->

[eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/FS_SystemFsTuzigoot.doc](http://www.in-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/FS_SystemFsTuzigoot.doc)

This document together with the Functional Specs describes the complete service architecture.

Occasionally, this document describes component interactions above layer 3 for clarifying the architecture/component interactions.

4 Network Architecture For Video Service

Network architecture greatly affects service deployment issues such as scalability, quality of service, security, and provisioning. It is therefore important to understand the network architecture's influence on multicast video service and identify aspects of network architecture that help (or adversely affect) multicast video service deployment. The following sections describe the various aspects of network architecture highlighting their pros and cons for video service.

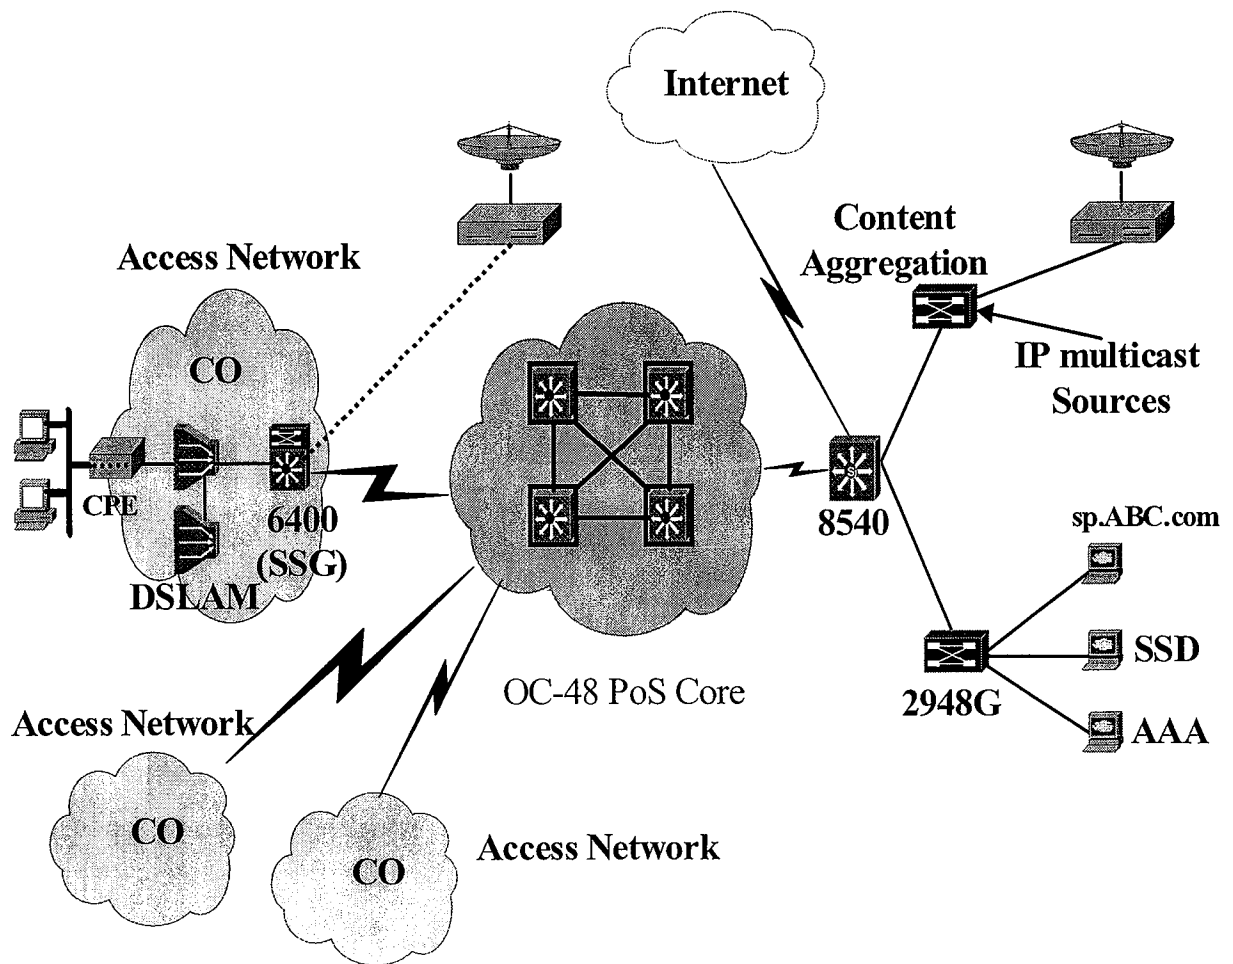


Figure 2 Network Architecture for Video Multicast Service⁴

⁴ A PC may interact with a Web server (sp.ABC.com) for obtaining information such as client software, or program guide. Additionally, AAA and Service Selection Dashboard servers will be needed for the video service. These interactions are described in "Tuzigoot System Functional Specifications" (ENG-41377).

The end-to-end network architecture can be logically divided into the Content Acquisition, Transport network (core), and Access network (please see).

A service provider can obtain video content by multiple means—via satellite (possibly from a content provider), stored content from a video server or from a live camera. The input streams, converted to IP multicast streams if necessary, are aggregated and transported to the Access PoPs via the core network. Although the network architecture is independent of the video compression technology, it is expected that the content is converted into MPEG-2 IP multicast stream (other compression schemes such as MPEG-1, MPEG-4, Clearband can be used as well).

The video streams are multicast over the core all the way to the access network. Unlike the unicast scenario, only one copy of a multicast packet is transported (in the core) to a Central Office (access network) regardless of the number of subscribers accessing the stream in that CO. The router nearest to the subscriber in the access network dynamically (on demand from subscribers) replicates the multicast stream, and forwards the stream to subscribers, per industry standard IGMP protocol. Routers in the transport network and the access routers (in COs) must, therefore, be multicast enabled to support the video service.

A video channel's maximum bandwidth can not exceed the last mile bandwidth in the access network. Bandwidths ranging from 256 Kbps to 1.5 Kbps are typical in current xDSL based access networks, although higher bandwidths are also being deployed. In xDSL networks, the maximum last mile bandwidth that can be provisioned for a subscriber depends on the subscriber's distance from the central office.

The access network consists of the xDSL modem (CPE) at the subscriber premise, DSLAMs, and the access router. At the subscriber premise, one or more PCs are connected via a 10/100 Ethernet LAN to the CPE. The corresponding xDSL modem in the Central Office (ATU-R) is integrated into the DSLAM (Cisco 61xx, 62xx) that aggregates subscriber traffic and forwards the traffic (over OC-x) to the access router (Cisco 6400). The access router is the last multicast router that replicates down stream multicast video packets and forwards them to individual subscriber CPEs.

For video multicast service, the access network is very important in terms of scalability and performance. This document focuses on the suitability of the different types of prevalent access architectures, PPP over ATM, PPP over Ethernet, Integrated Routing & Bridging (IRB), and Route Bridge Encapsulation (RBE)⁵.

5 Content Acquisition Architecture

A service provider may receive video streams from content providers or may generate video streams locally⁶. A video stream is typically compressed in MPEG2 (MPEG1/MPEG4) format and is transported as IP multicast streams.

Video streams are typically aggregated at a central site, and transported to the PoPs using the transport network. Alternatively, video streams can be transported to each PoP via a content provider's private satellite network, bypassing the core. The latter option is very attractive since it conserves core bandwidth and makes the video quality independent of the

⁵ Cisco Proprietary. Formerly referred to as half-bridging. Available in IOS 12.0(5)DC.

⁶ A service provider may provide local video channels, or insert local advertisements into video channels obtained from content providers. This is done by temporarily switching the regular multicast stream with one carrying the local advertisement.

Quality of Service (QoS) configurations of the core network, although it adds additional equipment and increases management overhead.

5.1 Content Format

Video content is typically encoded in MPEG 2 and is transported as IP multicast streams. Since a client software in the PC (e.g., media player) decodes the video, any video-encoding scheme (MPEG1⁷, MPEG2⁸, MPEG4⁹, Clearband) can be supported as long as the client software can decode it.

1. For performance reasons the IP packet size used for video streaming should be made as large as feasible, since larger IP packet size leads to less number of packets for the same amount of data, and hence reduces routing overheads. Also ideally the packet size should fit evenly into ATM cells (payload size of 48 bytes), i.e., it should be a multiple of 48 bytes.
2. It will be probably necessary to multicast a single content (e.g., CNN) at different speeds (e.g., 300 Kbps and 600 Kbps) to suit different capacity xDSL lines.

5.2 Content Feed

The IP multicast streams (MPEG-2 over IP) can be obtained from multiple sources:

- From satellite. In case of CoolCast service, a StarGuide receiver receives the data from satellite and acts as the IP multicast source. While this has the advantage of video streams being unaffected by transport network's QoS issues, it distributes the management of the multicast source to the COs.

Note on Coolcast content Acquisition: Since Coolcast content is transmitted via a private satellite network, we can place Satellite receivers (provided by CollCast) at the PoP/CO (where 6400s are located) instead of taking a single feed in a central site and using the core for transport. This prevents unnecessary loading of the core network.

- From a video server (e.g., Clearband Server), that encodes/multicasts video camera or VCR output in real time, or multicasts pre-encoded stored content.
- From a Webcam with optional equipment as necessary to multicast MPEG IP stream.

⁷ MPEG1 is optimized to fit into a bandwidth of 1.5Mbits per second (the data rate of uncompressed audio CDs and DATs). Typically MPEG1 is compressed (hardware) in non-real time and decompressed in real time (hardware or software).

⁸ MPEG2 is intended for higher quality video for products like the "set top box". MPEG2 runs at data rates typically between 4 to 9Mbps (although lower/higher rates are possible), and is suitable for MPEG-2 delivering Cable and HDTV broadcasts. MPEG-2 is backward compatible with MPEG1. Both hardware (chip sets) and software (e.g., Clearband) solutions exist for MPEG-2 compression/decompression. Additionally, there is a specification for MPEG2 adaptation over ATM AAL5.

⁹ MPEG4 is a low bit rate compression algorithm intended for 64Kbps connections. MPEG4 is targeted at a wide-range of applications including mobile audio and visual applications and electronic newspaper sources.

5.3 Content Aggregation

Before transporting the video streams through the core network, they may need to be aggregated as shown in the following diagram:




Figure 3 Content Aggregation¹⁰

The content aggregation network could consist of 2948G¹¹ switches, aggregating Ethernet data streams from content servers, and consolidating them onto gigabit Ethernet uplinks to an 8540¹² router. This router(s) may set the IP precedence bits before submitting the streams to the switch core, and additionally may serve as an ingress point for Internet traffic. IP precedence set here at the content aggregation is used subsequently in the transport network and the access networks to provide appropriate QoS.

6 Transport Architecture

The core transport network needs to support the aggregated bandwidth of each CO, as well as high bandwidth multicast video content (if video is transported over the core). An OC-48, or an OC-12 core may be suitable depending on the aggregated traffic volume. Compared to IP over ATM over SONET, PoS (Packet over SONET) provides 25- to 30-percent gain in efficiency (by eliminating ATM cell header, IP over ATM encapsulation, and segmentation

¹⁰ Although content aggregation can be done at multiple sites, typically a service provider prefers to do it at a central site.

¹¹ A 2948G can be used as an aggregation switch. It has the capability to EtherChannel its' 2-Gigabit ports. This provides an aggregate throughput of 4Gbps per switch. This is enough bandwidth to support 4000 1Mbps video streams. The 2948G has a 24Gbps non-blocking switch fabric supporting 48 10/100 and 2 uplink Gigabit ports. The non-blocking aspect of the switch means that traffic will never get backed up because of switch congestion. This is a cost-effective way of aggregating multiple 100 Mbps streams into 1 pipe.

¹² Cisco 8540 is ideal for layer 3 routing into a OC-48 PoS core (other Catalyst series switches may be used for aggregation for lower bandwidth requirements). The 8540 will essentially be used as an IP aggregation device into the core of the network, aggregating the video stream traffic as well as Internet, Extranet, and management network traffic. The reason for using the 8540 over another platform is due to the port densities for Gigabit Ethernet and the availability of an OC-48 PoS interface.

and re-assembly [SAR] functionality), and so is preferable. Any existing transport architecture can be used to transport multicast video streams¹³, as long as it has the required bandwidth, and supports appropriate multicast routing protocol and QoS.

Since it is simple to estimate the aggregate bandwidth required for video multicast streams (they being fixed), the total core bandwidth requirement can be estimated by adding the aggregated bandwidths of multicast video to the existing core traffic volume.

6.1 Transport Network Components

The ideal components for the transport network would be Cisco 12000 GSR's connected via OC-48 PoS. As mentioned above, PoS provides inherent efficiencies over ATM. The GSR's would reside in each major central office. Each CO that doesn't have a GSR will have to be connected to CO's that do via an ATM VC.




Figure 4 Transport Architecture (OC-48)

For an OC-48 core, a GSR 12012 would be an ideal choice:

- 2 x 1-port Gigabit Ethernet Line card (only 1 GSR needs this for the management network feed)
- 1-port OC-12 line card for the GSR connected to the content acquisition source
- 2 x 1-port OC-48c PoS line card

Note: depending on core bandwidth, other Cisco products may be used.

6.1.1 Transport Core Network Redundancy

Each node in the core should, at a minimum, be connected to two other nodes. This would provide redundant paths to each node in the case of failure. The management network ideally should have two independent paths into the core, but 2 paths into a single

¹³ A layer 2 core (ATM switch based) can be used for transporting IP multicast streams to the PoPs; however it will not be able to take advantage of IP multicasting (at layer 3).

node will suffice. If GSRs are used, for complete hardware redundancy, each component installed GSR should be redundant.

6.1.2 Multicast Routing Recommendation

IP Multicast uses multicast routing protocols for transport through a network of routers. Several of these routing protocols have been proposed through the standards bodies over the years (DVMRP, PIM, CBT, MOSPF). PIM V2 (Protocol Independent Multicast) routing protocol is on the standards track in the IETF, and has emerged as a de-facto standard and the protocol of choice for several reasons:

- Widely implemented in production environments
- More scalable than DVMRP and MOSPF
- Unlike DVMRP and MOSPF, PIM works with any existing unicast routing protocol (e.g., OSPF, EIGRP) implemented in the core
- CBT is new and does not have adequate acceptance as yet

PIM ensures that only one copy of a multicast packet is forwarded on each branch of the distribution tree, and a loop-free multicast topology is created for a specific multicast tree. PIMv2 takes advantage of intelligent unicast path determination mechanisms offered by protocols such as OSPF and EIGRP when constructing these distribution trees. .

PIM V2 can operate in dense mode, sparse mode, or sparse-dense mode. In dense mode, a router assumes that all other routers want to forward multicast packets for a group. The packet is forwarded from the router down the SPT (Shortest Path Tree, the fewest number of links between a sender and receiver pair) to its receivers requesting the group multicast packet. The flood and prune behavior of Dense Mode, in which an upstream router queries PIMv2 neighbors at intervals for receivers of group multicast streams, places heavy WAN overhead on PIM neighboring links.

In sparse mode, a router assumes that other routers do not want to forward multicast packets for a group, unless there is an explicit request for the traffic. In Sparse-Mode a multicast source sends multicast traffic to a Rendezvous Point (RP) which is the root of the shared tree. The root distributes the content down the shared distribution tree toward the receivers.

Sparse-Dense Mode is generally recommended by Cisco, and is the recommendation in the Tuzigoot network architecture. Sparse-Dense mode is preferred because it enables a hybrid SM-DM environment that allows some heavily accessed channels to be configured in dense mode and others in sparse mode. When running Dense-Mode and Sparse-Mode in independent regions to support densely and sparsely populated receivers, the two regions will not communicate.

Using PIM Sparse-Dense mode will allow for hierarchical IP Multicast deployment support for both densely and sparsely populated areas, allowing communication between both environments, with the support for both source and shared distribution trees. This will provide the existence of other IP Multicast services and not prohibit those services from communicating across networked hierarchical boundaries

Sparse Dense mode also supports Auto RP feature, a Cisco proprietary standalone protocol used to distribute group to RP mapping without much configuration. Another auto RP advantage is that it allows configuring multiple routers as RP for a multicast group. Auto RP enabled routers periodically negotiate among themselves and select one router to the acting RP for a multicast group. Failure of a RP will lead to automatic selection of one of the back ups, adding robustness to multicasting.

Using PIM V2 Sparse-Dense mode allows for the use of a bootstrap router (BSR) to quickly distribute RP information. PIM Version 1, together with the Auto-RP feature, can

perform the same tasks as the PIM Version 2 Boot Strap Router (BSR). Both versions can be used simultaneously to support a mixed RP mapping environment. However, please note that Auto-RP feature is available on both PIM versions, while BSR, which is part of PIM v2 standard, is available with only PIM v2.

Consult your Cisco PIM documentation for more discussion about PIMv2.

http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/113t/113t_2/pimv2.htm#xtocid24521

Recommendations:

1. PIM Sparse-dense mode is recommended because it offers better flexibility.
2. Different video channels can be in sparse mode, or in dense mode depending on user viewing patterns.
3. The router closest to multicast source would be the rendezvous point (RP), if video is transported over core.
4. The NRP can be RP if CO receives multicast stream from satellite. However, for a multi-NRP 6400, a separate router may receive the multicast streams from the satellite receiver and multicast them to all the NRPs. If so, then that router should be the RP.
5. If it is required to reduce IGMP join time for some heavily accessed multicast streams, the edge router can be configured to statically joining the specific multicast groups. This essentially runs the network in a PIM dense mode for the multicast group, but latency between stream selection is reduced.

7 Quality of Service (QoS)

Multicast video service is quite sensitive to packet loss/delay since these adversely affect the video quality. Minimizing such loss/delay in a network calls for assigning higher priority to video traffic over other (less sensitive) traffic. This enables the network to drop (or delay) lower priority data favoring higher priority traffic, thus mitigating or avoiding traffic congestion. A network's QoS configuration controls such preferential treatment to specific traffic classes; hence QoS plays a very crucial role in multicast video service deployment.

The service provider should classify all the deployed services per their business criticality and loss/delay sensitivity and accordingly prioritize their data for QoS. Such clearly defined QoS policy helps in its implementation. Typically, multicast video streams will be classified as the highest priority traffic (however, voice over IP traffic, if present, will likely require higher priority).

Although traffic can be classified and controlled in various ways (source, destination, protocol, port), the layer 3 IP precedence (TOS bits in the IP header) provides a convenient way to assign priority to traffic.

7.1 Content Aggregation: Quality of Service (QoS)

Recommendation: Committed Access Rate (CAR) in content aggregation stage.

For QoS to be effective, each packet needs to have its IP precedence bits set at the edge of the network, before the packet enters the core network. Otherwise, packets may get delayed before reaching the core.

In IOS there are two ways to set the IP precedence bit, Policy Routing¹⁴ and Committed Access Rate (CAR).. It is recommended to use CAR to set IP precedence before the multicast stream enters the transport network. CAR also allows packets to be flagged and/or dropped if they maintain or exceed static bandwidth configurations CAR is implemented with a simple configuration, and its rate limiting capability can be optionally used to control the amount of bandwidth used for IP multicast streams from a content source.

7.2 *Transport: Quality of Service (QoS)*

Recommendation: Weighted Random Early Detection (WRED) in the core (if an IP based core is used).

The core network may also get traffic congestion since it carries traffic of all services and to all PoPs. Congestion avoidance mechanisms that selectively drop lower priority packets to avoid congestion are preferable in the high-bandwidth core. Note that if multicast video streams are directly received from satellites at the PoPs (bypassing the core), any existing QoS policy of the core will not be influenced by multicast video service deployment.

If congestion avoidance in the core is required, the service provider may consider deploying **Weighted Random Early Detection (WRED)**. WRED, being a congestion *avoidance* mechanism, actually tries to mitigate congestion by selectively dropping packets before congestion occurs. It drops lower priority traffic, favoring higher priority traffic (based on IP precedence). The effectiveness of WRED will depend on the traffic mix and volume, and the service provider's QoS traffic classification. WRED is highly effective in environments where there are high amounts of bandwidth (since it takes longer for maximum bandwidth utilization to occur), and is very useful with protocols that intelligently slow down transmission when packet drops are detected (i.e., TCP).

If a layer 2 core (ATM) is used to transport IP multicast streams, the virtual channels should be provisioned with appropriate Class of Service (VBR or CBR). IP Precedence should either be automatically mapped to ATM CoS (if such feature is available on the platforms), or such mappings are manually configured.

Note: There are two different methods for controlling QoS in the core; congestion management and congestion avoidance. Queuing algorithms such as Weighted Fair Queuing (WFQ), Priority Queuing (PQ), and Custom Queuing (CQ) are used to manage congestion situations based on each packet's precedence. Packets with a higher precedence value are given priority in the output queues of the interfaces configured for queuing. Queuing is most effective in low speed networks since it tries to control congestion after it has already happened.

7.3 *Access: Quality of Service (QoS)*

¹⁴ Policy Routing allows you to not only set IP precedence, but to route packets using different links than those specified in the routing table. These alternate links could be faster or have different guarantees than the normal routed path. However, it does not have rate limiting capability.

Since the ADSL bandwidth to individual subscribers is small and this is where congestion is most likely to happen, the way we control congestion needs to change from preventative to controlling. WRED does not help us once congestion occurs. Controlling congestion as we discussed before is obtained by making intelligent queuing decisions. Instead of just pumping packets in or out as they arrive, packets with higher priority need to be scheduled first and packets with lower precedence are either moved back in the queue or dropped.

WFQ is recommended as the QoS mechanism in the access network. PQ makes queuing decisions based strictly on the precedence of an IP packet. CQ is like PQ except each queue is manually configured and it handles packets in a round-robin fashion. WFQ essentially breaks up the usable bandwidth based on the current precedence of the packets currently queued. This allows for a statistically more balanced queue.

Summary of end-to end QoS:

- Use CAR to set the IP precedence at content aggregation (e.g., on 8540 in content aggregation stage). If the video content is directly received at a CO, set IP precedence before it enters the access network (i.e., the 6400). Note that we are using CAR only for setting IP precedence, but not for setting bandwidth limits.
- Use WRED in the core (see section 7 Quality of Service (QoS), above)
- WFQ in access router (6400 NRP) is recommended, if available.

8 Access Architecture

Access network architecture plays a crucial role in multicast service deployment in terms of scalability, security, quality of service, and subscriber provisioning. Scalability is an important consideration due to the higher bandwidth requirement of video, and some access architectures scale better than others for multicast service.

Security is an important aspect of any consumer service such as the multicast video service. Certain access network architectures that are susceptible to IP hijacking, for example, are not suitable for residential service deployment.

The Quality of Service (QoS) configuration in the access network is another important consideration for providing quality video in the face of access network congestion. Without higher priority for video traffic, video loss or jitter will result.

Finally, subscriber-provisioning must be manageable with reduced complexity for large-scale service deployment to be viable.

The access architecture should support the following categories of services (list not exhaustive):

- Broadcast channels (basic and premium): IP multicast streams from 300 Kbps up to 1.5 Mbps.
- At least premium broadcast channels require conditional access.
- PPV: special case of broadcast channels. Needs event based conditional access.
- WebCam type applications. Generally lower bandwidth IP multicast stream. Restricted access needed; but a few seconds for authentication is acceptable.
- Internet access
- VPN services
- Proxy services

Access Bandwidth for video Service:

The multicast video service will be delivered over xDSL based access networks that provide adequate bandwidth for the video contents described above. The service requires high downstream bandwidth for the video stream, while minimal upstream bandwidth is needed. This fits well with the bandwidth asymmetry (high downstream & low upstream) available to most xDSL subscribers today (e.g., ADSL)¹⁵. Unused upstream/downstream bandwidth can be used by other services being accessed simultaneously. Currently, typical ADSL consumer links have 384 to 640kbps bandwidth downstream, although higher bandwidth (e.g., 1.5 - 2mbps) are also being offered. Since higher the bandwidth, better the video quality. The bandwidth of the access network available to a subscriber determines the video quality in a DSL network. It is possible to get quarter (or half) screen video of acceptable quality at the lower bandwidths (e.g., 300Kbps), while full screen video of higher quality would be possible at higher bandwidths.

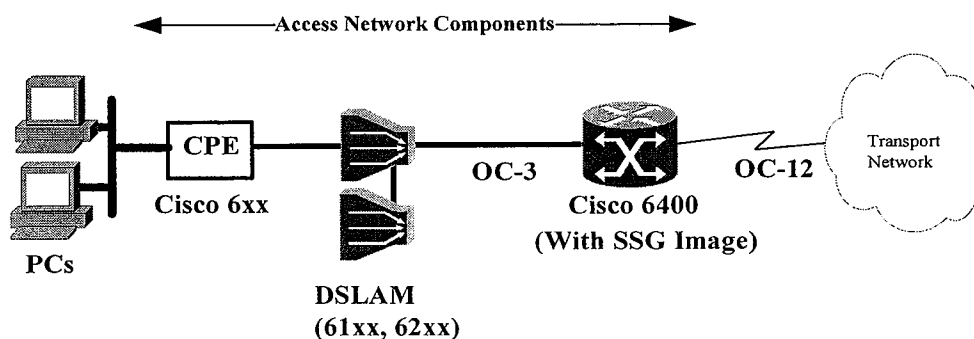


Figure 5 Access Architecture

Access Components and Configurations:

At the customer premise, one or more PCs are connected via an Ethernet LAN¹⁶ with a Cisco 6xx CPE (which can be configured to operate in RFC 1483 bridge, or PPP routing mode). On its WAN or ATM side, the CPE is connected to a DSLAM (Cisco 6100, 6130, 62xx)¹⁷. The CPE is configured to bind data to a well-known PVC (e.g., 0,1). Each CPE is physically connected to a separate DSLAM port. A unique PVC is provisioned per

¹⁵ Some xDSL (e.g., SDSL) deployments offer symmetric upstream/downstream bandwidth, where most upstream bandwidth will be available to other services. WebCams connected to residential xDSL links can benefit from such higher upstream bandwidth.

¹⁶ In future, the architecture will also support a PC with the recently announced INTEL's DSL modem card. A PC with such a card will be connected to a DSLAM directly without needing a NIC. The card's functionality is similar to Cisco's 605 modem (now discontinued) and supports drivers for both 1483 bridging and PPP routing. Therefore it will support RBE and PPPoA. The card is expected to support PPPoE as well (to be investigated).

¹⁷ The DSLAM can be a Cisco 6100/6130/62XX series, configured with up to 32 line cards for main subscriber access shelf, an OC-3 NI card for ATM connection to the SSG/NRP network, and with subtended modules and nodes if required.

subscriber CPE at the DSLAM's upstream side, and bound to the subscriber's VPI/VCI on the DSLAM's downstream side.

The 6400¹⁸, a combination of an ATM switch (NSP) and routers (NRP), aggregates high capacity feeds from the DSLAMs. On the downstream side, the NRP acts as the last multicast router for video streams; it replicates multicast packets, and forwards them downstream. The Service Selection Gateway (SSG) feature of 6400 NRP allows a user to connect simultaneously to multiple destinations, and is also required for supporting authenticated access to multicast streams. Additionally NRP-SSG¹⁹ has been modified as part of the Tuzigoot project to authenticate IGMP join requests based on the user's AAA profile, for the purpose of demonstrating such authentication.

- **PC:** a PC is expected to have pre-loaded (or dynamically acquired) client software needed to display the multicast stream. The PC IP address assignment mechanism varies per access technology used, and will be described in subsequent sections²⁰.
- **CPE:** the CPE is essentially an ATU-R, with additional functionality (it can act as a DHCP server and a router). It supports a downstream Ethernet interface that connects to an Ethernet LAN; its upstream ADSL interface is connected to telephone line to the CO, where it is connected to a DSLAM port. The CPE can be configured to operate in RFC-1483 bridge mode (supporting PPP over Ethernet, or route bridge encapsulation), or in RFC-1483 routing mode (supporting PPP over ATM).

Different Cisco CPE models are intended for residential and business subscribers. The target CPE for the residential subscribers is the 6xx because of its low COGS.

- **DSLAM:** A DSLAM aggregates and forward the xDSL traffic over ATM to a Universal Access Concentrator (Cisco 6400). Cisco offers a several DSLAMS (6100, 6130, 6260) with a wide range of features.
- **6400 UAC:** In the access network, DSLAM's are connected to the Node Line Cards(NLC) in 6400 chassis via OC-3. The 6400 Node Switch Processor (NSP) in turns maps incoming PVC's from the DSLAM to individual NRP-SSG cards. The NSP is basically an LS1010 ATM switch. Each NRP-SSG card is an independent 7200 class router running IOS version 12.05 DC that can support 2048 VC's. Each 6400 NSP can support 32,000 VC's. The 6400 chassis supports full hardware redundancy called EHSA, Extended High System Availability. Each blade is fully redundant with its odd numbered partner (i.e. 1 and 2, 3 and 4, etc...). In a fully redundant configuration the 6400 will have 4 active OC-3 ports or 1 active OC-12 port, 4 inactive OC-3 ports or 1 inactive OC-12 port, 3 active NRP-SSG's, 3 standby NRP-SSG's, 1 active NSP, and 1 standby NSP. This means a total of 6,166 VC's maximum per chassis.

The 6400 is connected to the transport network via OC-12 link.

The multicast stream may be transported to a 6400 from the core network, or from a local satellite receiver. For multi-NRP 6400, it required to have another router (e.g. 7200) between the receiver and the 6400 to multicast the incoming video streams to the NRPs.

¹⁸ More information about Cisco CPEs, DSLAMS, and 6400 can be found at <http://www.cisco.com/warp/public/cc/cisco/mkt/access>

¹⁹ SSG image in NRP is required for authenticating IGMP join requests. If such authentication is not required, then SSG is not mandatory for IP multicast service, although it may be necessary for other services being deployed over the network

²⁰ If any PC needs permanent IP address (e.g., to support HTTP server), then it should be either statically assigned, or the appropriate DHCP/AAA server should be configured accordingly..

Multicast Support in Access Network:

The 6400 UAC acts as the last multicast router, i.e., on getting a downstream multicast packet, it replicates and forwards the replicated packets to members of the associated multicast group. It runs multicast routing protocol PIM Sparse-Dense²¹ for interaction with upstream routers, and uses IGMP to interact with the subscriber PCs.

Multicast Support via Tunnels

When a user accesses his/her VPN service, the 6400 UAC (with SSG functionality) establishes a L2TP tunnel from the 6400 NRP to a corporation's home gateway. The user may access (via the tunnel) any IP multicast stream offered by the corporate network. Please note that the access network, in this case, simply carries the multicast IP packets in the tunnel, without in any way taking part in multicast routing. It is expected that multicasting is enabled in the corporations network. By utilizing this, a corporation can deliver a multicast video stream over a tunnel to a remote xDSL user.

Different Access Architectures:

The access components can be configured differently to implement the following access architectures:

- **Integrated Routing & Bridging (IRB):** allows a router to act as both bridge and router on the same interface. A large number of xDSL deployments today use IRB due to its simplicity and familiarity.
- **PPP over ATM (PPPoA):** PPPoA is the access architecture endorsed by the ATM Forum. A large number of deployments exist today.
- **PPP over Ethernet (PPPoE):** relatively new access architecture with limited deployments. It is an IETF standard, and requires each PC to have pre-installed PPPoE client software. Current Windows platform do not have this built-in software, but third party client software is available. Many feel that the requirement for pre-installed client software and ensuing support issues have hampered PPPoE's acceptance.
- **Route Bridge Encapsulation (RBE):** this is an improvement over IRB that leads to scalable multicast service deployment. This is a new technology from Cisco that specifically addresses the multicast scalability and security issues of bridged networks.

8.1 General Multicast System Flows

This section describes the general system flows common to RBE, PPPoA, and PPPoE access architectures. Flows specific to individual architectures are described in subsequent sections.

²¹ Static Multicast Mapping at 6400: if required, a 6400 NRP can be configured permanently join specific multicast groups to reduce channel switching latency.

8.1.1 User/Service Authentication

User and Service authentication takes place above layer 3 and is outside the scope of this document; however since user log-in/authentication is prerequisite for the rest of the flows, it is briefly described here for improving clarity.

- After the PC is booted and acquires an IP address (if not statically assigned), the subscriber would open his/her web browser and access the web-based Service Selection Dashboard (SSD) via HTTP.
- SSD presents an Account Logon page to the user, asking the user to enter his/her user-id/password. The user enters user-id/password. This user response is sent to an AAA server in the service provider network.
- After a successful logon, the end user is presented (by SSD) with a list of services available to him/her.
- The user selects a service. The user selection is again authenticated by SSD/SSG.
- On successful service authentication, SSD turns the service icon green. The user now can access the service. For a multicast video service, the user's browser is redirected to the Web page associated with the service. The Web page allows the user to select a video channel and access the video content within a Web page.

Please note that the user is expected to subscribe to one or more multicast services, each providing a set of video channels. The user will be allowed to access a video multicast channel only after successfully logging in to a multicast service that includes the channel.

Please refer to Tuzigoot Systems Functional Specs (ENG-41377) for details on the user interface and multicast service description:

http://wwwwin-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/FS_SystemFsTuzigoot.doc

8.1.2 PC issues a IGMP Join Request

Prerequisites:

1. For successful processing of IGMP join request, it is required that:
 - The subscriber has successfully completed user log-in using SSD
 - The subscriber has successfully completed multicast service log-in for a service that multicasts the requested stream.

This lets the NRP-SSG get the required user/service profile information to authenticate subsequent user IGMP requests.

2. PC client software is aware of "Channel Name to multicast group address" mappings, so that it can request for the correct IGMP group in response to a user selecting a channel (see Tuzigoot System Functional Specs, ENG-41377).

System Flow:

- PC sends an IGMP request to 6400 NRP, requesting to join a specific multicast group.
- IOS forwards the packet to SSG module. SSG module performs user access authentication (per user's AAA profile). If the user can access the multicast group, SSG module forwards the request back to IOS.
- On getting the packet IOS adds the user's sub-interface to the multicast routing table for the requested group. In this case, downstream multicast packets destined for the group will be forwarded to the sub-interface.

- If the user is not authorized for the multicast group, SSG module detects this from the user's AAA profile, and drops the packet. Since the IGMP request is dropped, the user will not get the IP multicast stream.

Note: The exact SSD and user interaction is described in the Tuzigoot Systems Functional Specs (ENG-41377).

8.1.3 PC issues a IGMP Leave Request

- PC issues IGMP leave to 6400 NRP
- If "immediate leave" is configured (see note below), 6400 NRP-SSG removes the user's sub-interface from the multicast table for the requested group. The PC stops getting the multicast stream.
- If "immediate leave" is not configured, NRP-SSG sends "Last Member Query Count" Group-specific IGMP queries every "Last member query Interval" to the multicast group being left, seeking to know if any one is still listening to the group. If any PC sends an IGMP report indicating that it is still receiving the multicast group, NRP-SSG does not remove the sub-interface from the multicast group table. Otherwise, it removes the sub-interface from the table, ending transmission of multicast stream to the sub-interface.

Default value for Last Member Query Count =2, Last member query Interval= 10 (1 second). These values are hardcoded in IOS.

Note: "Immediate leave" can be configured using a hidden IOS 12.0(5)DC command. Without this command, issuing an IGMP leave immediately followed by a Join (e.g., channel switching) may make a PC receive both multicast streams for a few seconds, possibly resulting in xDSL link overload (see 10.1 Channel switching latency and IOS modification/configuration for details)

"Immediate Leave" configuration solves the problem, but introduces a new one in situations where more than one PC are connected to a CPE and receive the same multicast stream simultaneously. If one PC leaves the multicast group, NRP immediately stops forwarding the multicast stream to the user's sub-interface, i.e., to both the PCs.

This may not be a problem if only one PC is expected to be connected to a CPE.

8.1.4 User issues a IGMP leave for a channel not authorized to him/her

IOS ignores the request since it does not find the 'user to multicast-group' association in the multicast table.

8.1.5 User issues multiple IGMP joins (for multiple channels) without leaving channels

The user will start getting all multicast streams authorized to him/her. The DSLAM-CPE line bandwidth may be insufficient, resulting in packet loss.

8.1.6 Hacker Sends garbage IP multicast streams to a multicast group address

It is possible for multiple sources to multicast IP streams to a single IP multicast group. A router will however accept a multicast stream input from a single interface (per Reverse Path Forwarding)²²..

If a subscriber (within the walled garden) arbitrarily sends IP multicast packets to a multicast group of a video service, the packets will be dropped by NRP. In this case, the packets from the subscriber will reach the NRP at an interface different from the one via which it normally receives content for the multicast stream, and the same router (NRP-SSG) can not accept packets for a multicast group on more than one interface per reverse path forwarding (RPF).

²² ACLs can be used, as an additional level of security, to accept multicast packets from known sources.

8.2 Access Architecture with Integrated Routing & Bridging (IRB)

The access network can use IRB, RBE, PPPoA, and PPPoE. This and the following sections describe the suitability of each of these modes of operation.

IRB is not recommended for IP multicast service due to the following reasons:

Bandwidth Issue:

- IRB forwards a multicast packet to all members of a bridge group. So a downstream multicast packet will be forwarded to all CPEs in the bridge group. PC-DSLAM link gets many times more (useless) traffic; DSL link capacity is overloaded, leading to video disruption.
- IRB also broadcasts an ARP request to all members of bridge group, affecting scalability.

Security Issue

- ARP spoofing possible, since NRP sends an ARP request to all members of a bridge group.
- Even if IOS subscriber-bridging feature is used to control ARP flooding, layer 2 to layer 3 address mapping is still learned via ARP. IP hijack is still possible.

Note 1: The above security issue is not all that serious for enterprise (the system administrator will catch an offender), but it is critical for consumer service delivery.

Note 2: Putting one CPE per bridge group will eliminate the above problems, but it leads to scalability issues due to IOS limit of 256 bridged groups.

8.3 Access Architecture with Route Bridged Encapsulation

8.3.1 Introduction

Routed Bridge Encapsulation (RBE) is a new feature in IOS 12.0(5) DC that alleviates the security and bandwidth issues associated with plain bridging (or IRB). It is suitable for providing video multicast service over xDSL networks.

RBE differs from IRB in handling upstream/downstream packets, and eliminates the limitations of IRB, as stated below:




Figure 6: Access Network Architecture (RBE)

Upstream packet:

IP packets originating from the customer premise arrive at a routed bridge interface as RFC 1483 Ethernet frames. However, instead of bridging the Ethernet or 802.3 frame, the NRP ignores the bridge header, and routes the packet to an outbound interface (packets of other protocols are processed as usual).

Downstream multicast packet:

A down stream multicast packet is forwarded to only the members of the multicast groups, instead of the whole bridge group (as happens in IRB). So a CPE does not get multicast packets intended for other CPE's, as happens in IRB (thus optimizing bandwidth usage).

ARP Request:

Unlike IRB, RBE does not flood an ARP request to the whole bridge group; instead an ARP request is sent to the specific interface for the IP address.

- Since each subinterface is associated with one residential subscriber, the possibility of ARP spoofing is now limited to a single household (having multiple PCs);
- Improves scalability

Per CPE Subnet:

RBE uses per CPE subnet. Since each subinterface has unique subnet, faking the IP address (from another subscriber's subnet) in an ARP reply is detected by NRP; it would drop the packet generating "wrong cable" error.

RBE Characteristics:

- CPE is configured for RFC1483 bridging.
- The subscriber's interface on the NRP is configured to perform "routed-bridging" (IOS allows Bridging, PPPoE, and routed bridging to co-exist on the same subinterface).
- A bridge group (or BVI) need not be defined.
- RBE is supported for IP (not for other protocols such as IPX, AppleTalk).
- one NRP subinterface is associated with a single subscriber (single PVC).
- each subinterface is assigned a unique subnet (a /30 or /29 subnet can be used that will allow 2 or 6 IP addresses per subnet. The subscribers premise can not have more than the allowed number of PCs).

8.3.2 Protocol Stack

PC	CPE		DSLAM		6400 UAC
IP					IP
Ethernet	Ethernet	1483 bridging LLC/SNAP			1483 Bridging (LLC/SNAP)
		AAL5			ALL5
		ATM	ATM	ATM	ATM
10/100BaseT	10/100BaseT	DSL	DSL	OC-x	OC-x

8.3.3 Address assignment options for the PCs and CPE

- In RBE, the CPE (in 1483-bridged mode) does not have an IP address and plays no role in assigning IP addresses to the connected PCs nor in any address translation (PAT).
- The PCs are normally configured to obtain IP addresses from a DHCP server (unless addresses are manually configured).
- The NRP terminating the CPE's PVC is configured as the DHCP relay agent, and forwards the DHCP packets from the PC to a DHCP server. The NRP inserts the address of the receiving subinterface in the giaddr field in the DHCPREQUEST message. Therefore the DHCP server will assign an address from a matching scope. It is expected that distinct IP pools be defined such that each subinterface belongs to a different subnet.

Public and Private IP Space:

PCs can be assigned private or public IPs. For private IP, NRP-SSG is configured to perform NAT²³.

If private IP addresses are used, IP loss due to subnetting is not an important issue. For public IP, the loss of addresses due to subinterface ip assignment can be mitigated by making each sub-interface unnumbered and adding a static route for each CPE. This however needs statically assigning IP addresses to PCs (not by DHCP).

The following table summarizes the various RBE address assignment options with their pros and cons:

	Private/Public address for PC	Static IP address, or DHCP used	Subinterface has IP address or is ip un-numbered?	Pros/Cons
1	Private	DHCP	With ip-address	Best option, if private IPs can be used (loss of private IP does not matter)
2	Private	DHCP	Ip-unnumbered	Not feasible, since DHCP will assign addresses from a single scope to all PCs (not unique subnet per CPE, as needed by RBE)
3	Private	Static	With IP-address	Worse than #1, due to manual PC IP configuration. If DHCP can't be used then this may be considered. No static route needed.
4	Private	Static	IP un-numbered	More configuration intensive than #3, since in addition, static routes need to be provisioned.
5	Public	DHCP	With IP address	Best option for public IP. Subinterface assigned an IP address No static routes needed. Easiest provisioning in RBE, but some IP loss due to subnetting.
6	Public	DHCP	IP un-numbered	Not feasible (same reason as #2)
7	Public	Static	With IP-address	Configuration intensive. IP loss due to per CPE subnet. No static routes. Manual IP configuration for PC. Uses more IP than #8. But easier to provision (no static route).
8	Public	Static	ip un-numbered	Configuration intensive. Subinterfaces do not use ip, but IP loss due to per CPE subnet. Static routes needed. Manual IP configuration for PC.

²³ Note that multicast addresses are precluded from NAT. However, if private IP space is used, NAT will be needed to access other non-multicast services, or to access some server in the Internet as a prerequisite for accessing multicast streams (e.g., CoolCast Web server).

Recommendation: #5 or # 1 are preferable, in public and private IP address spaces respectively.

Note1: if private addresses are used, then NRP-SSG needs to perform NAT. IPsec can't be used if NAT is used.

Note2: In Private IP space, IP loss is not an important factor; so such loss is not highlighted always in the above table.

8.3.4 Configuration Guidelines (6400 NRP)

- Each CPE PVC is terminated at an NRP-SSG subinterface; a subinterface has at most one PVC.
- Subinterface is configured for route-bridge (i.e., encapsulation aal5snap, atm route-bridge ip)
- Configure QoS parameters for downstream traffic (e.g., for UBR)
Note: without such QoS parameter, 6400 NRP sends full wire speed data downstream leading to connection drops.
- A subinterface is associated with a unique IP subnet; this can be done by:
 - Assigning an IP address from a distinct subnet to each subinterface, and using a DHCP server to assign CPE IP addresses from these subnets (NRP acts as DHCP relay agent), or
 - Use IP un-numbered subinterfaces, with static routes defined per CPE.

Note: in the first approach, if a /29 subnet is used, there are a total of 6 usable IP addresses; the subinterface is assigned one, leaving 5 for the customer premise PCs. The DHCP server is configured one IP pool per subscriber subnet. In the above /29 subnet example, the DHCP will have an IP pool consisting of at most 5 addresses.

8.3.5 Pros & Cons of Route Bridge Encapsulation mode for providing IP multicast based services

Pros:

1. Good option to migrate from IRB to support multicast service, since the CPE operates in standard 1483 bridge mode.
2. Simple CPE configuration (compared to PPPoA); no per user configuration necessary.
3. CPE vendor independence.
4. Performance:
 - Upstream better than IRB (since IOS directly routes the packet, without doing bridging classification of the packet).
 - Downstream performance will be better due to:
 - No multicast flooding to whole bridge group
 - No ARP flooding
5. Scalability:
 - IDB use is same as PPPoA (1 per ADSL loop)
 - Being stateless, RBE may scale better than PPPoA/PPPoE
6. Support for WFQ at 6400 NRP:

Weighted Fair Queuing is supported on a RBE subinterface on 6400 NRP. So if there is network congestion on the subscriber's subinterface (possibly due to simultaneous access to multiple services), WFQ shall drop the subscriber's lower priority traffic favoring his/her multicast video traffic.

PPP virtual access interfaces (e.g., PPPoA and PPPoE) support FIFO only. Please note that, for commonly accessed services such as Web access (HTTP), e-mail, ftp that are TCP based, FIFO may not cause appreciable degradation of video quality, since TCP traffic slows down on detecting packet drops, allowing the UDP traffic (multicast video) to use sufficient bandwidth.

Cons:

1. If DHCP is used, each subinterface (on 6400 NRP terminating a subscriber PVC) needs IP, to help DHCP find correct scope²⁴. Not an issue if private IP is used.
2. Alternatively, if NRP subinterface is ip unnumbered (and no DHCP is used), then one static route per subinterface is needed; this is somewhat configuration intensive.
Workaround: use script to generate NRP configuration
3. Some IP loss due to per CPE subnetting. If availability of sufficient public IP is an issue, consider private IP with SSG NAT.

²⁴ With the currently planned "DHCP unnumbered" feature, DHCP will work with unnumbered sub-interfaces.

8.4 Access Architecture with PPP Over ATM

8.4.1 Introduction

The PPP over ATM access architecture is suitable for video multicast service. This section describes PPPoA access architecture, its pros & cons for multicast service and applicable address assignment options,

The PPP-over-ATM feature (RFC 2364) enables the NRP-SSG to terminate multiple remote Point-to-Point Protocol (PPP) connections. These PPP connections are typically received from each subscriber CPE.

A logical interface, known as a virtual access interface, associates each PPP connection to an ATM permanent virtual circuit (PVC). You can create this logical interface using the `atm pvc` command. This configuration allows the PPP protocol to terminate at the router ATM interface as if received from a typical PPP serial interface. Each PPP connection is encapsulated in a separate ATM PVC, which acts as the physical medium over which PPP frames are transported.

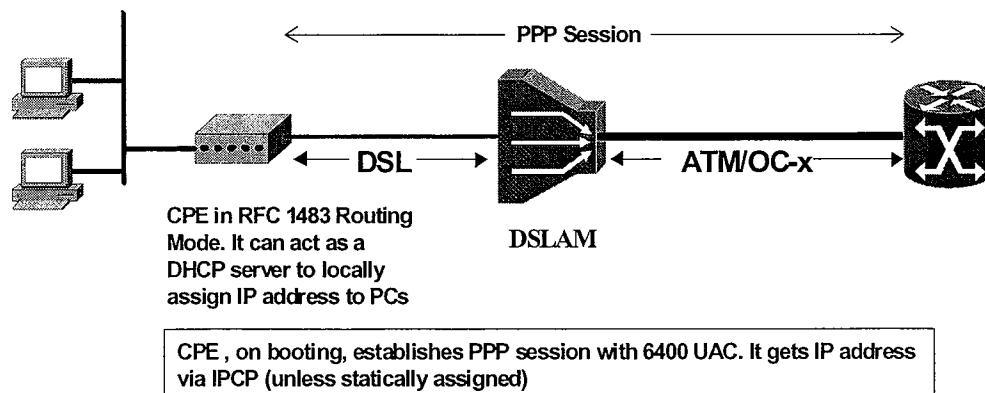


Figure 7: Access Network Architecture (PPPoA)

The virtual access interface for each PVC obtains its configuration from a virtual interface template (virtual template) when the PVC is created. All PPP parameters are managed within the virtual template configuration. Multiple virtual access interfaces can spawn from a single virtual template; hence, multiple PVCs can use a single virtual template.

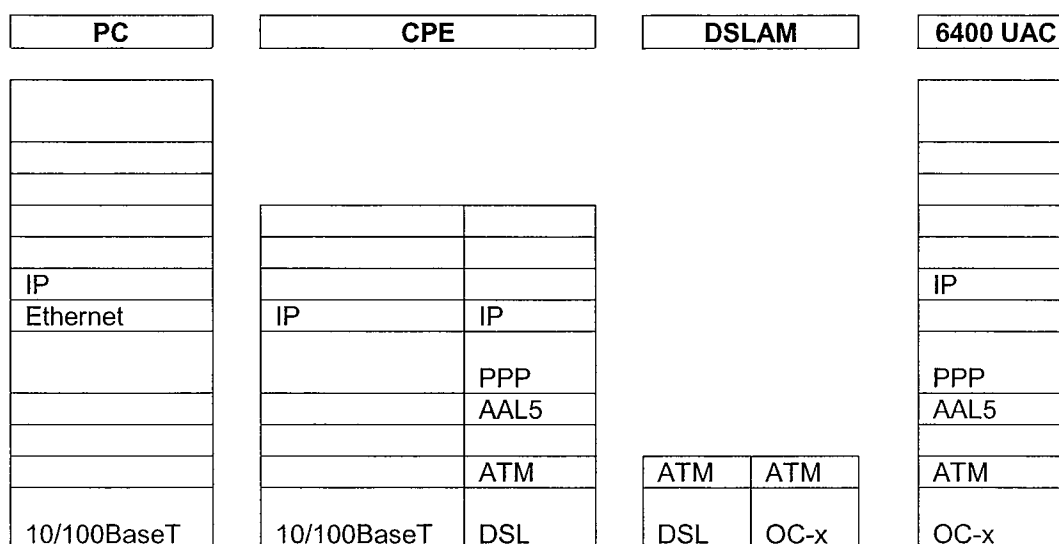
PPPoA requires per-subscriber CPE configuration with user name and password for PPP authentication and an IP address pool. To avoid such per-subscriber CPE configuration, service providers often configure all CPEs with identical user names and passwords, and live with the fact that no real PPP authentication of a CPE is possible. Cisco provides Cisco IOS features in release 12.0(5)DC to address this. These features enable mass CPE configuration in PPPoA environments while allowing individual CPE authentication. They are:

- VPI/VCI-based authentication: allows RADIUS authentication of individual CPEs by identifying CPE via its associated VC rather than a username, enabling

individual authentication of CPE provisioned with a default user name/password see Appendix C: VPI/VCI Based CPE Authentication).

- IPCP subnet feature: avoids manual configuration of CPE DHCP pools. This feature enables CPE to automatically configure its local IP pool based on the subnet mask it receives during PPP negotiation. See Appendix D: IPCP subnet feature:

8.4.2 PPPoA Protocol Stack in Access Network



8.4.3 Address assignment options for the PCs and CPE

- CPE IP address can be typically assigned or via PPP/IPCP
- PC's can be assigned static addresses if the IP address and subnet mask returned to the CPE is always the same.
- If the IP address and subnet mask assigned to the CPE are different each time the PPP session is renegotiated, PC's can be configured to retrieve IP address information via DHCP from the 675.
- If NAT is enabled on the CPE, private addresses can be assigned via DHCP since the only address visible to the outside world is that assigned to the CPE (in this case, the CPEs do not need individual subnets).
- If CPE NAT is not used, then each CPE belongs to a separate subnet (like in route-bridge case), i.e., there are as many subnets as there are subscribers.

8.4.4 Configuration Guidelines(PPPoA)

6400 UAC:

- Create a virtual-template

- Configure the virtual-template for PIM dense mode or the appropriate PIM method
- Create either a point-to-point or multipoint ATM sub-interface
- Configure the PVC of the subscriber (i.e. pvc 5/54)
- In PVC configuration mode, configure the encapsulation type for AAL5MUX PPP and the virtual-template created in the first step (i.e. encapsulation aal5mux ppp Virtual-Template1)
- Multiple PVC's can be configured under one interface. One way to simplify configuring multiple PVC's for the same attributes is to use VC class mappings.
- In global config mode create a VC-CLASS (i.e. vc-class atm <name>) and configure it for any attributes normally configured under an ATM VC.
- Under the ATM sub-interface for multiple PVC's, configure it to use the class from the previous step (i.e. class <name>).
- Now every PVC under that sub-interface inherits whatever is configured under the global VC-CLASS.
- Subscriber PVC's do not need to be assigned to individual sub-interfaces. PPPoA VC's can be grouped under a single interface so that default ATM classes can be globally mapped to the sub-interface.
- The ATM Class of Service (CoS) and Virtual-Template can be configured under the "vc-class" global configuration option and then referenced under the ATM sub-interfaces that need those values. This allows for simplification of the configuration process by not having to repeat common values or options.

8.4.5 Pros & Cons of PPPoA for IP multicast based Services

Pros	Cons
Since the CPE establishes a PPP session, it can (if implemented in future) negotiate and find out the least common denominator features between itself and the access router. So if a CPE is upgraded to have enhanced capability (e.g., encryption) is installed, it can turn it off or on depending on whether the NRP supports it or not. This will not be possible in RBE, or PPPoE, which run CPE in bridged mode.	CPE configuration is not as easy as it is for RBE and PPPoA, and typically includes a subscriber's name and password, and address pool for the local subnet (for the subscriber PCs). Unless VPI/VCI based Authentication and IPCP subnet features are used, this would preclude CPE mass configuration.
Security via PPPoA authentication (of CPE) is helpful in non-SSG environment	For video multicasting, SSG is likely to be used for IGMP authentication; PPPoA authentication will be redundant due to SSG's user/service authentication
CPE can NAT subscriber traffic. Address space is conserved.	Application's that embed IP information in the payload can not be NAT'ed without special handling. Example: Netmeeting, CUSeeMe, and ICQ. If CPE NAT is used, then SSG gets only one (translated) IP address for all the PCs behind the CPE. Hence it can only form one host object for all the PCs, thereby restricting all PCs to have only one SSG session simultaneously.

PPP session traffic can be accounted for (if SSG is used, then such accounting is redundant due to SSG's accounting functionality)	6400 NRP supports only FIFO queuing on PPP virtual access interfaces
VPI/VCI based CPE authentication can be used enabling CPE configuration on a per user basis (please see 13 Appendix C: VPI/VCI Based CPE Authentication)	VPI/VCI based authentication needs specific AAA server that can support the feature (e.g., Cisco's Access Registrar).

8.5 Access Architecture with PPP over Ethernet

8.5.1 Introduction

PPP over Ethernet (PPPoE) is suitable for providing video multicast service. PPPoE provides the ability to connect a network of hosts over a CPE configured for RFC1483 Bridging to the 6400 NRP-SSG. With this model, each host utilizes its own PPP stack and the user is presented with a familiar user interface. To provide a point-to-point connection over Ethernet, each PPP session residing on a host must learn the Ethernet address of the remote peer such as the 6400, as well as establish a unique session identifier. PPPoE includes a discovery protocol that provides this (see RFC 2516 for details).

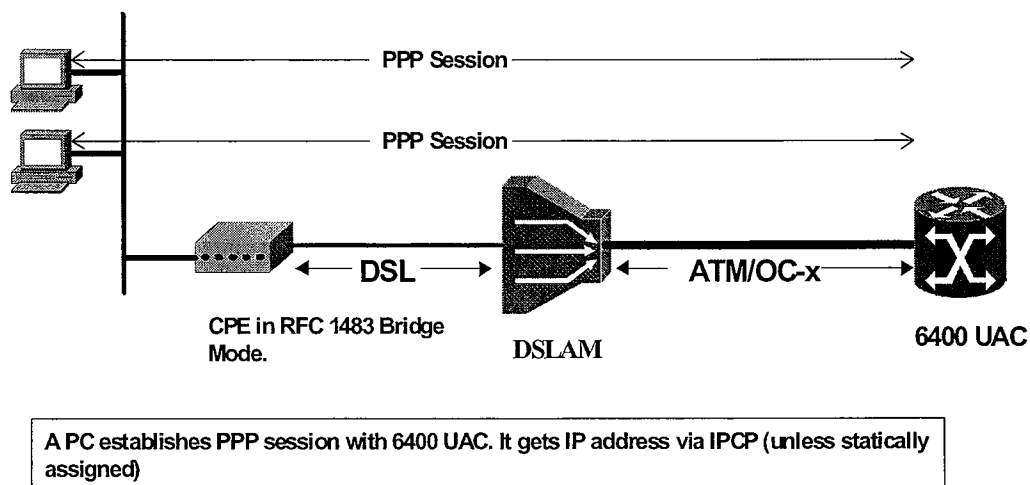


Figure 8: Access Network Architecture (PPPoE)

8.5.2 Protocol Stack

PC	CPE		DSLAM		6400 UAC
IP					IP
PPP					PPP
PPPoE					PPPoE
Ethernet	Ethernet	1483 bridging (LLC/SNAP)			1483 Bridging (LLC/SNAP)
		AAL5			AAL5
		ATM	ATM	ATM	ATM
10BaseT	10BaseT	DSL	DSL	OC-x	OC-x

8.5.3 Address assignment options for the PCs and CPE

- CPE runs in 1483 bridged mode; does not have IP.
- A PC behind a CPE establishes a PPPoE session that terminates at 6400 NRP; PCs get IP via PPP/IPCP.
- Alternatively PCs may be manually configured with IP address.

8.5.4 Configuration for 6400/6100/CPE/PC

- Configure a virtual-template interface
- Enable VPDN (i.e. vpdn enable)
- Configure a VPDN group to accept-dialing for the PPPoE protocol and specify the virtual-template previously configured.
- Create an ATM point-to-point of multipoint sub-interface
- Under the ATM sub-interface configure the subscriber PVC and specify the protocol as PPPoE
- If multiple VC's are to use the same attributes, a VC-CLASS can be assigned. Refer to the configuration instructions in the PPPoA section.
- CPEs need not be under different subnets.

8.5.5 Pros & Cons of PPPoE for providing IP multicast based services

Pros	Cons
Very suitable for offering familiar PPP interface to subscribers while maintaining a RFC 1483 bridged CPE environment.	Requires installation of 3 rd party software on each system at the customer premises (only supports Win32 and MacOS; WIN2000 is stated to include this). Support issues for the third party software may limit deployment.
Simple CPE configuration. CPE is configured for bridged mode.	6400 NRP supports only FIFO queuing on PPP virtual access interfaces.
Each user at a subscriber's home can be authenticated individually (even without SSG). However, this is redundant if SSG is used.	
Does not need per CPE subnet, so equally well suited to public/private IP environments.	

8.6 Recommendation (RBE, PPPoA, PPPoE)

Although any of the RBE, PPPoE, or PPPoA architectures can be used for deploying scalable multicast video services, the following advantages of RBE should be considered when choosing one.

- In case of access network congestion, RBE may offer better QoS via Weighted Fair Queuing, which is not available to PPPoA or PPPoE
- Unlike PPPoA, RBE and PPPoE need simple (RFC 1483) CPE configuration that makes multi-vendor interoperability possible

- Unlike PPPoE, RBE requires no third party client software installation or support in subscriber PCs. PPPoE requires each workstation behind a CPE to install a third party PPPoE stack. In mass deployments, supporting the third party client software becomes a critical issue

8.7 Scalability/sizing: (for RBE, PPPoA, PPPoE cases)

The number of subscribers that can be supported by the access network for multicast video service depends on several factors:

– Multicast Stream Characteristics:

- Video Stream Bandwidth: more subscribers can be supported by a NRP for a smaller bandwidth stream than for a higher bandwidth stream (e.g., ~250 Vs. ~160 maximum for 500K, and 300Kbps streams respectively, as described subsequently).
- Multicast Packet size: Multicast packet size should be close to the access network MTU, for better routing performance²⁵.

– NRP –DSLAM link traffic:

The OC-3 link between the 6400 and the DSLAM carries the REPLICATED multicast packets, e.g., if 200 subscribers are online and each watch one of 5 (300Kbps) channels being broadcast, the link carries $200 \times 300\text{Kbps} = 60\text{Mbps}$, not $5 \times 300 = 1.5\text{ Mbps}$ (packets are in fact replicated from the 6400 NRP to the NSP).

8.7.1 Calculating the total available downstream capacity of 6400:

The NRP-NSP back-plane (155 Mbps) carries both upstream and downstream traffic.

Available bandwidth on the OC-3 link minus protocol overhead ~ 125 Mbps.

Upstream traffic for multicast service is extremely low (e.g., IGMP reports from the CPE) compared to the downstream traffic, and can be ignored.

So the total downstream bandwidth available ~125 Mbps for OC-3.

This bandwidth carries the replicated multicast packets.

8.7.2 Number of Subscribers that can be supported

The 6400 chassis has 8 slots to be shared by NRP cards and line cards. Each NRP card takes one slot. Each OC-12 line card (with one OC-12 port) occupies one slot. An OC-3 line card (with 2xOC-3 ports) takes ½ slot (so one slot can have two OC-3 line cards with a total of 4 OC-3 ports).

²⁵ For example, Clearband packet size is 1410 bytes (configurable for Clearband server).




Figure 9: Sample Multi-NRP Cisco 6400 Configuration

The above figure shows a redundant 6400 configuration, having 6 NRPs (3 pairs of NRPs, each pair having a redundant NRP²⁶), 2 OC-3 cards (each with 2 OC-3 ports), and an OC-12 card on the trunk port for connecting to the core. Three OC-3 ports are used to connect to 3 DSLAMs (which can be subtended).

A non-redundant configuration could be 5 NRPs in 5 slots, a trunk OC-12 card in one slot, and 5 OC-3 ports (3xOC-3 line cards) taking one full slot and half of another.

A single NRP can deliver from 45 Mbps (worst case; full replication) to 105 Mbps (best case; no replication needed) of multicast output (as explained subsequently)²⁷. If the video streams are of 300K, then taking the middle ground of say 75 Mbps, gives us 220 concurrent video sessions.

The following table lists the min/max NRP performance, and the # supportable concurrent users with various stream bandwidths:

	Min.	Avg Case	Max. Case
NRP Performance	105 Mbps	75 Mbps	45 Mbps
# of video users (300 Kbps stream)	340	230	125
# of video users (500 Kbps stream)	210	150	90
# of video users (1Mbps stream)	105	75	45
# of video users (1.5 Mbps stream)	70	50	30

Min: when NRP is doing 100% replication; meaning all the downstream users want to watch the same channel causing the NRP to do full replication

²⁶ N x 1 redundancy (i.e., 1 redundant card NRP per N other NRPs) is slated to be available with NRP2 – this will increase the total routing throughput of the chassis.

²⁷ NRP2 will increase multicast throughput

Max: when each user is watching a unique and different channel. Here NRP acts as a mere forwarding engine (no replication)
Avg: (Min + Max)/2.

The following table states the number of supportable users per 6400 chassis assuming average replication case figures as above:

	Non-redundant (5 NRPs)	Redundant (3 NRPs)
# users (300 Kbps Stream) @ 230/NRP	1150	690
# users (500 Kbps Stream) @ 150/NRP	750	450
# users (1 Mbps Stream) @ 75/NRP	225	525
# users (1.5 Mbps Stream) @ 50/NRP	250	150

Please note that the number of video subscribers for subtended DSLAMs connected to a NRP should be limited (via subscriber provisioning) to what a NRP can support (see above).

9 SSG based Authentication for IP Multicast group

Currently (IOS 12.0), IGMP traffic is transparent to SSG; any user can join any IP multicast stream. This is suitable for deploying multicast video services based on single flat rate subscription to all provided channels on the network.

For the Tuzigoot project, SSG (Vulcan image in 6400) has been enhanced for IGMP join authentication (in a private branch), as follows (please see Tuzigoot Multicast Authentication Functional Specs ENG-44252 for details):

- Broadcast Channel (ip multicast group) to be configured as a SSG compatible 'service', or as part of a service.
- The AAA user profile will list the services accessible by a user.
- IOS, on receiving a PC's IGMP join request for a multicast IP address, will forward the packet to SSG.
- SSG will authenticate the request based on the user's AAA profile.
- If the user is not allowed to access the multicast group, SSG will drop the packet; else it forwards it to IOS resulting in the join to take effect.

The above changes to IOS enables a service provider to offer subscription based multicast video service. Although not productized yet, this feature is currently demonstrable in the lab.

Note1: SSG Image requirement for Multicast Authentication.

- 6400 NRP with SSG image is required for SSG based Authentication for IP Multicast groups.
- Further, if a 6400 has more than one NRP, the NRP that terminates the multicast subscriber's PVC must have the SSG image. Otherwise, the NRP terminating the PPP always honors all IGMP joins from users, without doing any authentication. It treats the other NRP (with SSG enabled), as another router and interacts with it via configured multicast routing protocols (e.g. PIM) to process IGMP Join requests.

Note: relaxing the above requirement (of SSG to be enabled on the NRP terminating the subscriber PPP sessions) would require some kind of IGMP proxy capability to be implemented on NRP.

Note2: SSG User/service Authentication/log-on procedure

Although system interactions at layer 4 and above are outside the scope of this document (these are provided in Tuzigoot Systems functional specs), for the sake of completeness, a brief description of the user log-on and service log-on procedure (when SSG, SSD, AAA server are used) follows:

After the PC is booted and acquires an IP address if not statically assigned, the subscriber would open their web browser and access the web-based Service Selection Dashboard (SSD) via HTTP. SSD presents an Account Logon page to the user. This account logon is sent to an AAA server in the service provider network. After a successful logon, the end user will be presented with a list of services available to him/her. For more details on the user interface, please refer to the Tuzigoot System Functional Specs at

[http://www-win-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/FS_SystemFsTuzigoot.doc](http://www.win-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/FS_SystemFsTuzigoot.doc)

Note3: SSG based authentication vs. Encryption based Authentication:

It is also possible to implement authenticated access to video channels by encrypting the video stream, and providing the subscriber PCs with appropriate keys for decryption.

- The authentication can be per user basis in SSG based authentication. So the user can be mobile. In encryption based scheme, the encrypted keys are distributed to a specific PC).
- In encryption based authentication, the content provider encrypts and hence controls access. In SSG based solution, the control is with the Service Provider.
- Changes in subscription to channels will lead to AAA user profile modification in SSG based scheme. SSG based authentication will be very suitable for offering premium channels. Supporting PPV services, where large number of user profile updates will happen in a short time window before the PPV event starts, will be limited to the transaction handling capacity of the AAA server (SSG based authentication), or key distribution server (encryption based authentication).

10 Technical Issues

10.1 Channel switching latency and IOS modification/configuration

When a subscriber issues an IGMP Leave for a multicast group followed by a join for another group (as would happen when he/she switches broadcast channels), it is possible that both streams could simultaneously be sent to a subscriber for a few seconds. This is due to the time it takes the IGMP leave message to be processed, as explained below:

On getting the IGMP leave request, NRP stops forwarding the multicast stream, but only after it queries and confirms that no other user is accessing the channel via the sub-interface. For confirming this, NRP-SSG sends "Last Member Query Count" Group-specific IGMP queries every "Last member query Interval" to the multicast group being left, seeking to know if any one is still listening to the group. NRP-SSG continues to forward the multicast stream until all these queries time out (or a PC responds saying it is listening to the stream).

An IGMP join request, unlike a leave request, is immediately honored. So if an IGMP leave is immediately followed by an IGMP join, for a few seconds (until NRP actually stops forwarding packets of the stream being left) the PC will get contents of both streams. If the ADSL line bandwidth is NOT sufficient for both channels, then packets will be dropped, leading to video loss.

Default value for Last Member Query Count =2, Last member query Interval= 10 (1 second). These values are hardcoded in IOS.

Workaround:

IOS 12.0(5)DC provides a hidden command that causes the NRP-SSG to stop forwarding the specified multicast stream to the user's sub-interface, without sending IGMP queries. The command is: "ip igmp immediate-leave group-list <acl>" in global or interface config mode, where <acl> is an access list that permits/denies the multicast groups to leave immediately.

This works well if one PC is connected to a CPE. If multiple PCs in a household watch the same channel simultaneously and one leaves the channel, the others stop getting the channel.

Note: This may not be a problem if only one PC is expected to be connected to a CPE.

IOS Enhancement Needed:

An alternate approach would be to modify IOS to let the "Last Member Query Count" and "Last member query Interval" configurable (hardcoded at present). Then they could be set to low values in NRP-SSG reducing the channel switching latency.

Applicable DDTS reports: CSKdk78345 on 12/28/98,

10.2 Enforcing Max Channel Limit

NRP can honor any number of IGMP joins from a subscriber, which may lead to xDSL link overload. There is a need to have a mechanism to restrict a subscriber's access to either a

fixed number of multicast streams or to allow as many as can be handled by his/her DSL link. This is an open issue at present.

Note: NRP-SSG uses the "ssg maxservice <#>" parameter to globally enforce the maximum number of services each subscriber can access simultaneously. This feature does not help, since it is in terms of number of services rather than multicast streams, and it is not on a per user basis.

10.3 IP QoS to ATM CoS

Currently mapping IP precedence to ATM Class-of-Service (CoS), does not exist in the 6400. The feature lets us direct IP traffic based on its precedence down VC's configured for different bit rates (i.e. VBR-rt vs. UBR). This allows us to take advantage of ATM's inherent capabilities of managing congestion and providing real-time data traffic. Adding this feature to the 6400's IOS feature set would provide even better QoS capabilities for the architecture described in this document. This feature can be supported when per vc-queuing is supported in hardware.

10.4 Fair DSLAM subtending

When DSLAMs are subtended, any DSLAM in the tree can utilize a disproportional amount of bandwidth to the 6400, starving the other DSLAMs. A feature to fairly subtend DSLAMs will greatly enhance the quality of video service by enabling service providers to subtend DSLAMs with the knowledge that bandwidth is fairly divided among the DSLAMs²⁸.

10.5 Smart DSLAM

The future DSLAM is planned to be IP aware to better manage traffic congestion and provide better quality of service. If DSLAM drops a layer 2 cell, being IP aware it knows that the entire IP packet is useless and so it discards the rest of the layer 2 cells of the IP packet. This is much better than the normal behavior of dropping consecutive cells, and thus invalidating all the IP packets they belong to²⁹.

10.6 Support WFQ for PPP virtual Access Interfaces

6400 NRP supports only FIFO queuing on PPP virtual access interfaces. WFQ, PQ, CQ (which are typically configured on the main interface), if configured are ignored for a PPP virtual access interface. This means that for PPPoA and PPPoE, WFQ (or CBWFQ) is not effective on subscribers' sub-interfaces on 6400 NRP.

Supporting WFQ for PPP virtual Access Interfaces (on NRP) will enable a service provider to further enhance the QoS capability of PPPoA and PPPoE subscribers.

Note: even without specific QoS configurations, video multicast stream (UDP packets) will indirectly get preference over TCP traffic (e.g., Web browsing, e-mail, ftp). This happens because TCP adjusts transmission speed when packets are dropped. Since UDP does not adjust, TCP ends up slowing down sufficiently favoring video traffic.

²⁸ This feature helps multicast and non-multicast services and is available in NI-2.

²⁹ Also in NI-2

10.7 Multicast Enabled DSLAM (Future)

In future, it would be ideal to make the DSLAM multicast enabled. This would make the architecture more scalable by:

- reducing IGMP processing overheads in NRP, and
- avoiding replicated multicast traffic across the NRP-DSLAM links

11 Appendix A: VLAN Based Architecture for IP Multicast Service

This is an alternate architecture based on VLANs and is suitable for providing higher bandwidth video services, although it can not support user authentication, service selection, and multicast authentication. If the primary goal is to provide higher bandwidth video service (2Mbps and above) and lack of user/service/multicast authentication is not of concern, then this architecture can be considered.

The end to end architecture and the Access Components are highlighted in Figure 10 End to End Architecture and Figure 11 Access Architecture respectively.

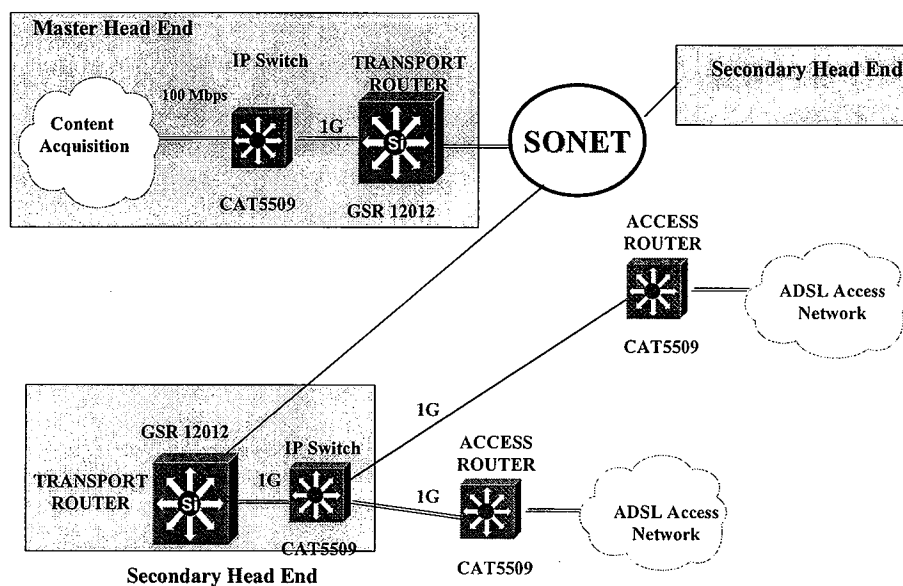


Figure 10 End to End Architecture

Content Aggregation:

At the master headend, video channels, i.e., IP multicast streams (on 100M Fast Ethernet) are aggregated by CAT5509 and forwarded to a GSR12012 via Gigabit Ethernet. The GSR 12012 connects to the SONET backbone via OC-12.

Transport Network:

The Transport network consists of GRS12012s in each PoP (OC-12 ATM or PoS). OC-48 may be used if traffic volume warrants.

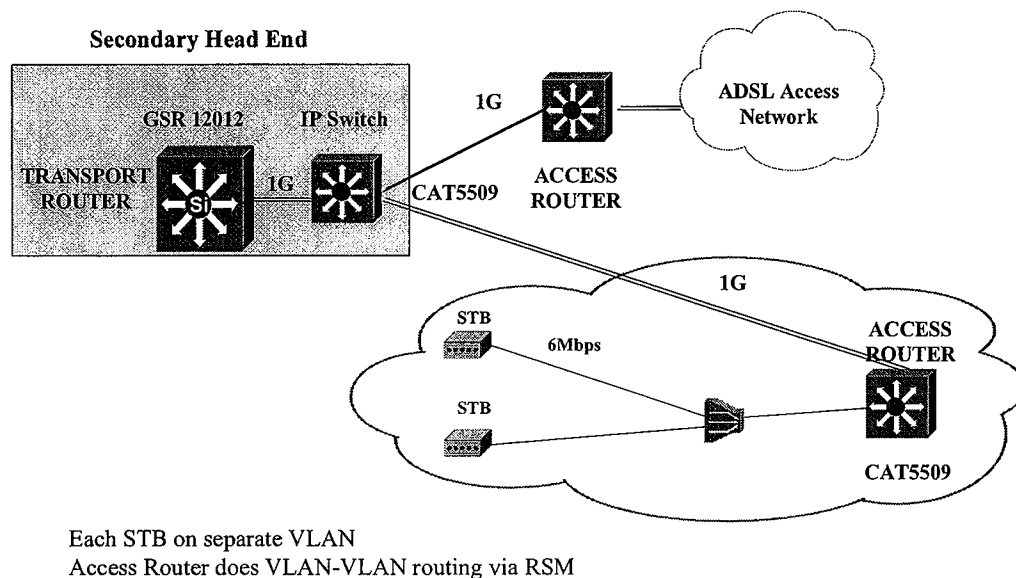


Figure 11 Access Architecture

Access Network:

- At the subscriber premises, a STB (or PC) is connected via 10/100 Ethernet to a CPE (ADSL Modem) that operates in 1483-bridged mode.
- The ADSL Modem (CPE) is connected via xDSL link to the DSLAM (downstream bandwidth should be sufficient for a broadcast channel)
- DSLAM to access router has OC-3 link (OC-12 preferable).
- The access router is a CAT5509 with RSM. It is the last point for multicast routing; it replicates multicast packets and forwards to individual members. CAT5k/6k platform has capability to replicate multicast packets in hardware (Multicast MLS feature), leading to excellent replication performance³⁰.
- Each subscriber has one PVC from the CPE to the access router (5509).
- On the Access router, each subscriber STB is on a separate VLAN. The Access Router RSM performs VLAN to VLAN routing.
- The Access Router is connected via Gigabit Ethernet to a IP switch (another 5509)
- The IP switch aggregates traffic into the GRS12012 via Gigabit Ethernet.

Although this architecture is a good choice for providing video service, it has the following limitations:

³⁰ A CAT 5000 with RSM, SUP-III and Netflow feature card can easily provide replication at higher than OC-12 speeds.

- No service selection via SSG
- No user authentication
- No multicast authentication (since no SSG)
- Scalability is limited, since the CAT5K – DSLAM link (OC-3) carries replicated video streams.

12 Appendix B: Filtering Multicast Traffic by ACLs

It is possible to control access to multicast IP streams by setting up per-user downloadable access control lists in SSG, but is not advisable. These ACLs can be defined in RADIUS user profile as cisco-avpair attributes. This is only true for PPP type subscribers. This would also preclude the use of SSG multicast authentication and be static for the life of the PPP session. Normally one ACL filter is used per IP multicast stream to be denied access (or allowed access); however if some multicast IP addresses to be filtered for a user happen to be in serial order then filtering those can be defined in a single ACL.

These attributes are defined per CPE in case of PPP over ATM. If PAT is not used, then ACLs can be used per user). When the PPP session is established between CPE and SSG a virtual interface is created and access control lists are applied.

Sample configuration:

To deny access to IP multicast group 239.217.63.147, define the following Cisco-avpair attribute in (Cisco Secure) the user profile:

```
ip:outacl#115 deny ip any 239.217.63.147 0.0.0.0
ip:outacl#120 permit any any
```

Although this approach can be used for regulating access to multicast IP streams, it is not scalable, and hence not recommended. For example, if an average user can not access 50 IP multicast streams (in practice, it may be more than this), then the user profile will have 50 ACLs (assuming non-contiguous group addresses). Even if an NRP supports only 300 users, the total number of ACLs will amount to 1500, which will severely affect the NRP performance.

Note: controlling up-stream multicast per subscriber:

Sometimes it is required to filter out upstream multicast traffic from certain subscribers (to prevent non-subscribers from multicasting video using Webcams, for example). This can be achieved by

- putting ACL on virtual access interface, or
 - RBE mode: put ACL filter on NRP and apply this filter to each ATM subinterface for RBE.
 - PPPoA/PPPoE mode: put ACL filter on NRP then applied to virtual-template interface
- put ACL filters on the Radius profile of the user (when user is logged in, SSG will apply the filters).

In ACL, you can put "permit igmp" but "deny all other class D addresses". The ip range 22.0.0.0 to 224.0.0.255 should not be blocked, because IGMP and other multicast routing protocol packets need to go through.

13 Appendix C: VPI/VCI Based CPE Authentication

This feature, to be available in 12.0(5) DC, helps to implement distinct RADIUS profile per CPE (in PPPoA environment), even when all the CPE are provisioned with a default user name/password.

In PPPoA, a CPE is assigned a user name and a password for RADIUS authentication. A NAP may however provision all CPEs with the same user name and password to reduce CPE provisioning efforts. This means all CPEs will be authenticated using a single RADIUS user profile; hence the RADIUS reply attributes (ip address, subnet, DNS server address) can not be configured on a per CPE basis.

VPI/VCI based authentication is a Cisco proprietary solution involving the NRP-SSG and needs an AAA server compatible with this feature. It exploits the fact that a CPE can be uniquely identified by the PVC connecting it to the 6400, and so authenticates a CPE by authenticating the VPI/VCI and some additional attributes associated with the PVC. This feature enables the NRP and the CPE to get CPE specific RADIUS configuration information, even though all CPEs are configured with a default user name and password. To use this feature, individual RADIUS profiles per CPE, identified by the CPE's PVC attributes, are created in a RADIUS server. So a CPE can be authenticated based on its individual profile.

The effects of CPE authentication using user-name/password vs. VPI/VCI are as follows:

NRP address assignment mechanism for CPE	Effect of user/password based authentication (one user profile for all CPEs)	Effect of VPI/VCI based authentication (one user profile per CPE)
External AAA server provides actual framed IP address.	Only a single user profile (so only one IP address) is used for all CPEs. Not feasible.	Individual (fixed) IP address (and DNSs address) can be configured.
External AAA server provides framed IP address indicating NRP to get IP address from an external DHCP server (RADIUS framed IP address must be configured to get IP address from a DHCP server)	Since one net mask is available, all CPEs are assigned addresses from the same scope.	DHCP can assign address (and DNS address based on different subnets assigned to CPEs.
NRP uses local pool.	Only a single (CPE) user profile (so only one IP address) is used for all CPEs. Not feasible	Not Applicable (VPI/VCI based authentication needs AAA server supporting the feature).

The VPI/VCI based authentication needs an external AAA server, and the following discussion assumes that one is used.

When a PPP session is initiated by a CPE, NRP authenticates the CPE by an external RADIUS server before terminating PPP. The authentication will be based on VPI/VCI of the CPE as reported on the 6400 ATM switch (NSP).

- On getting a PPP session request from a CPE, NRP-SSG forms a RADIUS packet and includes the user name and password in the packet. In addition, it inserts the following in the RADIUS packet:
 - NAS IP address = Management IP address of NSP
 - NAS Port address = slot+module+port+VPI/VCI, where
 - Slot = NLC slot where subscriber CPE connected.
 - Module = NLC module
 - Port = NLC port
 - VPI = VPI on the NSP for subscriber VC (not the interconnect vpi/vci of the NRP)
 - VCI = VCI on the NSP for subscriber VC (not the interconnect vpi/vci of the NRP)

4 bit slot + 1 bit module + 3 bit port	8 bit VPI	16 bit VCI
---	-----------	------------

- The RADIUS server is configured for VPI/VCI based authentication, ignores the user name and password attributes, and authenticates based on the NAS IP address and NAS port (subscriber provisioning includes creating a user profile accordingly; note that attribute 61 NAS-Port-Type is always specified as "virtual").
- The RADIUS reply items include the Framed-IP address, Framed Net-Subnet mask, and DNS server address (if configured).

Note: for backward compatibility, the AAA server recognizes standard user-id and password as belonging to older CPEs, and performs user-id/password based authentication (does not perform VPI/VCI based authentication).

For additional information please refer to "VPI-VCI Authenticated Routing AAA Feature Functional Spec" (ENG-33842) at

http://wwwin-eng.cisco.com/Eng/NUBU/xDSL/Solutions_Lab/VPI-VCI_Authenticated_Routing_AAA_Feature_FS.doc

and "6400 RADIUS VC Logging" (ENG- 33846) at

http://wwwin-eng.cisco.com/Eng/NUBU/xDSL/6400/6400_RADIUS_VC_Logging.txt.

For modification of Cisco Access Registrar (AAA server) to provide VPI/VCI based authentication, please see "Stray Cat System Functional Specification (ENG- 33757)" at:

http://wwwin-eng.cisco.com/Eng/NSMBU/StrayCat/Sys_Specs/straycatfunc.doc

14 Appendix D: IPCP subnet feature:

This feature to be available in IOS 12.0(5)DC, is useful in PPPoA environment for automatically configuring the CPE DHCP pool. It lets the CPE receive a subnet mask during PPP session establishment, so that it can accordingly configure its local pool for assigning IP addresses to its connected PCs.

With this feature, during setting up a PPP session a CPE gets authenticated and receives its subnet mask from the NRP (12.05 DC will return the subnet mask), in addition to the usual other attributes (CPE's IP address, and additionally DNS address as standard RADIUS reply attributes #135 and #136). The CPE and the PCs connected to the CPE belong to this subnet. Therefore the CPE can locally assign IP addresses from this subnet to the PCs. CPE provisioning includes defining (in the AAA profile) the correct subnet mask based on the number of PCs the CPE is expected to support.

On successful RADIUS authentication, the 6400 NRP uses the RADIUS reply (containing the framed IP address, framed net mask) to complete the IPCP negotiation with the CPE. It also installs routes in its routing tables to the CPE based on this IP address and subnet mask.

On the receipt of the IPCP negotiated IP address, the 675 CPE configures itself by doing the following:

- assigning the returned IP address to its local LAN Ethernet port
- configuring its wan0-0 port to run ip-unnumbered off of its LAN port address
- configuring the onboard DHCP server pool with a beginning IP address of "IPCP negotiated IP Address + 1" and the size of the pool being determined by the subnet mask. The pool will be configured with a default gateway of its LAN port address and DNS server IP address will be the one received as a part of the IPCP negotiation.

The PCs are DHCP enabled and when booted up they make a DHCP request in order to obtain an IP address. The 675's onboard DHCP server responds with an IP address, subnet mask, default gateway and DNS server information to the PC.

Note: IPCP Subnet is available in IOS 12.0(5)DC and higher

Note: CBOS 2.2 is required on the CPE for IPCP Subnet

15 Appendix E: Encryption Based Authentication for IP Multicast group

Encryption based authentication is an alternate to the SSG based authentication described above. Please note that such authentication is not explored in this document (it is not part of Tuzigoot); however its basic characteristics are provided below to help deduce its merits/demerits compared to SSG based authentication.

- Each multicast stream to be authenticated is encrypted with a key before being forwarded to the transport/access network by the content provider.
- The client application in the PC must have the key to decrypt the stream; otherwise the video stream display can't be intelligible.
- Keys for a multicast stream will be changed from time to time (e.g., every week for premium channels), or with PPV events.
- Keys should be distributed to the PCs from one/more key servers.
- Key distribution mechanism:
 - Ideally key distribution should be done at the POP level, to reduce latency.
 - Keys for a broadcaster's channel should be pushed to a PC from the POP level server.
 - PCs, on booting, must be able to retrieve its keys.
 - The authentication target is a PC, not the user (no user id is used). So the user can't use another PC to access the encrypted stream (this is similar to the existing cable STB model).

16 Appendix F: System Flows

16.1 System Flows in Route Bridge Encapsulation Architecture

16.1.1 CPE boots up/powered down:

CPE boots up:

- CPE trains for a few seconds with the ATU-R in DSLAM
- Traffic can now flow in both directions.

CPE is powered down:

If the CPE is powered down, the PCs stop getting IP multicast streams, and can't respond to the periodic IGMP queries from NRP-SSG; so NRP-SSG removes IGMP group memberships of the PCs (after the IGMP query interval).

- If the CPE is powered up before the next periodic IGMP query arrives, the PC will start getting the IP multicast streams it had joined. The PC will also respond to the IGMP queries, and continue to get multicast stream.

16.1.2 PC boots up/shuts down/powered down

PC boots up:

- PC gets IP address via DHCP and periodically renews lease
- NRP-SSG acts as the DHCP relay agent

PC is gracefully shut down:

- PC sends DHCPRELEASE; DHCP server receives it (via NRP-SSG) and releases the IP lease.

If NRP-SSG gets a DHCPRELEASE it processes it as any other IP packet and it does not destroy the Host and connection objects for the user. The user context in SSG still remains intact until the idle timeout/session time out expires.

- PC does NOT send IGMP Leave request(s) for any currently watched multicast stream(s). NRP-SSG continues to forward multicast streams to the PC for a while, until the PC does not respond to the periodic IGMP query from NRP-SSG (this interval is hard coded in IOS at present).

Note: the non-multicast based connections stay open; the user can reaccess them without requiring to enter his/her user-id/password.

PC loses power while multicast streams are being watched:

- This is exactly similar to above, except that the PC does not issue the DHCPRELEASE. The IP lease is not revoked in this case until it expires.

NOTE: if the PCs are configured not to use DHCP (IP addresses are manually configured), then the DHCP related interactions does not take place in the above flows.

16.1.3 Hacker Tries ARP spoofing

- In route-bridge mode, although NRP-SSG discovers the MAC address of the destination PC by ARP, it sends the ARP packet to only the destination sub-interface rather than all sub-interfaces of the bridge group (as happens in IRB). So a hacker can not listen to ARP requests for PCs other than his/her premises. This limits possibility of ARP spoofing to only the PCs behind a single CPE (a single household).
- If the user sends an ARP reply packet with correct MAC address, but some other user's IP address, the router will detect it and generate a "wrong cable" error. This is because each subinterface is configured with unique subnet.

16.1.4 User Configures the PC with arbitrary IP addresses

In RBE, each subinterface is on a separate subnet. If a PC is manually configured with the IP address of another user, the ARP reply from the PC will be detected to be from another subnet. NRP will generate a "wrong cable" error and the packet will be discarded.

Note that PCs within a single household will be under one subnet; hence they can possibly hijack one another's IP; this is not a serious issue since all PCs belong to one subscriber.

16.1.5 Subscriber to Subscriber Communication

- If the subscribers are behind a CPE (e.g., within a single household), then they can communicate on the local LAN.
- IP packets from a PC to another PC behind a different CPE are routed (not bridged)

16.2 System Flows in PPPoA

16.2.1 CPE boots up/Powered down

CPE Boots UP:

- The CPE initiates a PPP session to the access concentrator, the 6400 NRP-SSG. The NRP in turn validates the username and password via RADIUS by a configured AAA server.
- On successful RADIUS authentication, the 6400 NRP clones a Virtual-Access interface from the Virtual-Template associated with the PVC of the incoming PPP session and issues an IP address and subnet mask, via IPCP, to the CPE and creates an SSG host object for the PPP session. The capability to provide a subnet mask via IPCP is known as the "IPCP subnet feature" (please refer to Appendix D: IPCP subnet feature:) This allows the CPE to instruct its internal DHCP server to hand out addresses corresponding to the correct subnet that the CPE itself belongs to.
- The CPE configures its internal DHCP pool to consist of the address range:
Start address: CPE's IP address +1
End address: last valid address per returned subnet mask

Although IPCP feature is Cisco proprietary, it doesn't break compatibility with other vendor's PPP implementations. By setting up the CPE in this way, real addresses can be assigned to a subscriber and the use of NAT at the customer premises is eliminated.

Note: avoiding NAT(PAT) at customer premises is preferable since

- protocols that embed IP in the payload (e.g., H323, ICQ), and applications that expect fixed client port(s) can be supported.
- With NAT at customer premises, SSG sees only one IP address for all PCs behind the CPE, it creates only one host object for all of them; so all these PCs can access only the same set of services

CPE is Powered Down:

- PPP session between CPE and NRP is destroyed.
- SSG destroys the host object of CPE (along with its connection objects, if any)
- The users don't get any network traffic.

If the CPE is powered up immediately after being powered down, it re-establishes the PPP session with NRP-SSG.

Case 1: CPE NAT is used

The CPE host object is shared among all the users behind a CPE. On service log-on, a connection object is created associating the CPE host object with the service object. So destruction of CPE host object, on CPE shutdown, would also destroy all its connection objects. After power is restored to the CPE, it establishes a fresh PPP session, and a new CPE host object is created. But since the connection objects are lost, the user needs to perform service log-on again to access a service.

Case 2: CPE NAT is not used

In this case, SSG maintains separate host objects for the CPE and each user. A user's host object is associated with service objects via connection objects (CPE host

object is not). When the CPE is powered down, users can't send/get any network traffic. However, if the CPE boots up before the service time-out and idle time-out intervals, the users will be able to access services they are already logged into.

If the CPE remain down for more than service time-out period, the user's connection with the service is broken by SSG; the user needs to log-in to the service to be able to access it. If the CPE remain down for more than idle time-out period, the user's host object (and associated connection objects) is destroyed by SSG; user log-in as well as service log-in will be required to access services.

16.2.2 PC Boots Up/Powered Down

PC Boots UP:

- If configured with static IP address (subnet mask, and DNS address), the PC simply boots up and is ready for the user to log-in to the network.

If the PC is configured to receive its address via DHCP, it is assigned an IP address from the CPE, which acts as a DHCP server. After the PC receives its IP address as part of boot up, the user can log-in to the network.

PC gracefully shut down:

Similar to RBE.

PC looses power while multicast streams are being watched:

Similar to RBE.

16.2.3 Hacker Tries ARP spoofing

- Since CPEs have individual PPP sessions with NRP, ARP spoofing is not possible between PCs connected to different CPEs.
- Since PCs behind a CPE are connected by a LAN, ARP spoofing is possible among them; however, this is not an issue since the PCs belong to one household.

16.2.4 User Configures the PC with arbitrary IP addresses

User may configure the PC with a static IP, while the CPE is configured to give out IP via DHCP. If the user configures a legitimate address from the CPE's network space, then this isn't a problem. DHCP just doesn't get used (in the worse case the address may be same as another PC behind the same CPE; this will affect the other user within the household) If the user configures an address from outside the subnet, then CPE will not route the user's traffic.

16.3 System Flows in PPPoE

16.3.1 CPE boots up/Powered down

CPE boots up:

- CPE trains for a few seconds with the ATU-R in DSLAM
- Traffic can now flow in both directions.

Note: CPE is configured for RFC1483 bridging.

CPE is powered down:

If the CPE is powered down, the PCs stop getting IP multicast streams, and can not respond to the periodic IGMP queries from NRP-SSG; so NRP-SSG removes IGMP group memberships of the PCs behind the CPE (after the IGMP query interval).

Also, the PPP session between the user and the NRP-SSG is terminated. So after the CPE boots up again, the user has to re-log-in in order to access the services.

16.3.2 PC Boots Up/Powered Down

PC Boots Up:

The PC is configured not to use DHCP. After boot up, it is ready for the user to establish a PPPoE session.

PC gracefully shuts down:

The PPPoE session is terminated. SSG removes the host object of the PC, and all its connection objects. So the user need to establish the PPPoE session again and re-login to the network.

PC loses power while multicast streams are being watched:

Same as above.

16.3.3 PPPoE Session Establishment

To initiate a PPPoE session, a subscriber opens a Dial-Up Networking client (e.g., third party software from Routerware) and initiates the PPPoE dial session. A username and password field will be presented. This username must match that of a user account present in the 6400's RADIUS AAA server. After the subscriber inputs a valid username and password, the PPPoE stack goes through the following stages.

The client sends the PADI (PPPoE Active Discovery Initiation) packet with the DESTINATION_ADDR set to the broadcast address. The CODE field is set to 0x09 and the SESSION_ID MUST be set to 0x0000. The PADI packet MUST contain exactly one TAG of TAG_TYPE Service-Name, indicating the service the client is requesting (unused at present), and any number of other TAG types. An entire PADI packet (including the PPPoE header) MUST NOT exceed 1484 bytes so as to leave sufficient room for a relay agent to add a Relay-Session-Id TAG.

When the Access Concentrator receives a PADI that it can serve, it replies by sending a PADO (PPPoE Active Discovery Offer) packet. The DESTINATION_ADDR is the unicast address of the client that sent the PADI. The CODE field is set to 0x07 and the SESSION_ID MUST be set to 0x0000. The PADO packet MUST contain one AC-Name TAG containing the Access Concentrator's name, a Service-Name TAG identical to the one in the PADI, and any number of other Service-Name TAGs indicating other services that the Access Concentrator offers. If the Access Concentrator can not serve the PADI it MUST NOT respond with a PADO.

Since the PADI was broadcast, the client may (theoretically) receive more than one PADO from multiple access concentrators. The client looks through the PADO packets it receives and chooses one. The choice can be based on the AC-Name or the Services offered. The client then sends one PADR packet to the Access Concentrator that it has

chosen. The DESTINATION_ADDR field is set to the unicast Ethernet address of the Access Concentrator that sent the PADO. The CODE field is set to 0x19 and the SESSION_ID MUST be set to 0x0000. The PADR packet MUST contain exactly one TAG of TAG_TYPE Service-Name, indicating the service the client is requesting, and any number of other TAG types.

When the Access Concentrator receives a PADR packet, it prepares to begin a PPP session. It generates a unique SESSION_ID for the PPPoE session and replies to the client with a PADS packet. The DESTINATION_ADDR field is the unicast Ethernet address of the client that sent the PADR. The CODE field is set to 0x65 and the SESSION_ID MUST be set to the unique value generated for this PPPoE session. The PADS packet contains exactly one TAG of TAG_TYPE Service-Name, indicating the service under which Access Concentrator has accepted the PPPoE session, and any number of other TAG types. If the Access Concentrator does not like the Service-Name in the PADR, then it MUST reply with a PADS containing a TAG of TAG_TYPE Service-Name-Error (and any number of other TAG types). In this case the SESSION_ID MUST be set to 0x0000.

Once the PPPoE session begins, PPP data is sent as in any other PPP encapsulation. All Ethernet packets are unicast. The ETHER_TYPE field is set to 0x8864. The PPPoE CODE MUST be set to 0x00. The SESSION_ID MUST NOT change for that PPPoE session and MUST be the value assigned in the Discovery stage. The PPPoE payload contains a PPP frame. The frame begins with the PPP Protocol-ID (see RFC2516 for details).

LCP and IPCP interactions take place after the PPPoE session establishment.

Note: another way PPPoE can be utilized is using PTA. This allows a subscriber to have their traffic aggregated solely to a remote network. The process for authentication is the same as above except the IP address assigned to the PPPoE stack is returned by a remote RADIUS AAA server. This effectively makes a subscriber part of the remote network. This feature, however, is not directly useful for accessing IP multicast services.

16.3.4 Subscriber to Subscriber Communication

- If the subscribers are behind a CPE (e.g., within a single household), then they can communicate on the local LAN.
- IP packets from a PC to another PC behind a different CPE are routed by the 6400 NRP (except when they share the same NRP subinterface).

16.3.5 Hacker Tries ARP spoofing

Since PCs have individual PPP sessions with NRP, ARP spoofing is not possible between PCs.

16.3.6 User Configures the PC with arbitrary IP addresses

For the PPPoE session, any statically configured IP address is ignored.

17 Appendix G: NRP Configuration Matrix for PPPoE

Configure GLOBAL and VPDN-Group COMMANDS (Enable Multicast routing)	VPDN Must Be Enabled <i>vpdn enable</i> vpdn-group <number> (config-vpdn) accept dialin pppoe virtual-template <number> ip multicast-routing
Configure ATM Interface and/or Sub-Interface	interface atm slot/0.subinterface- number multipoint
Configure PVCs on ATM Interface or Subinterface	pvc <vpi/vci> ubr <bit rate>
Configure ATM Encapsulation Method & Protocol	encapsulation aal5snap protocol pppoe
Configure Virtual-template	interface virtual-template <number> ip mtu 1492
	peer default ip address pool <word> (the word is optional, example PPPoE)
Configure Global IP Local Pool	ip local pool PPPoE n.n.n.n n.n.n.n

Enable Multicast PIM on Interface	ip pim sparse-dense-mode Note: ensure you have turned on IP Multicast Routing.
AAA NRP LOCAL PPP Authentication	aaa new-model aaa authentication ppp default local aaa authorization network default local username <name> password <password>

Configure NRP for AAA RADIUS Server Authentication If using Server for Authentication	Radius-server host n.n.n.n auth-port 1645 acct-port 1646 radius-server key <key password> (key secret)
Enable PPP Authentication IPCP Negotiation	PPPoE Client Initiated (Example, WinPoet)
Enable DHCP Server	
Enable NAT (if using private IP)	

18 Appendix H: NRP Configuration PPPoA,

Configure GLOBAL and VPDN-Group COMMANDS (Enable Multicast routing)	ip multicast-routing
Configure ATM Interface and/or Sub-Interface	interface atm0/0/0 or (subinterface atm0/0/0.1 optional) Note: If IP is assigned to interface enable ip pim sparse-dense-mode
Configure PVCs on ATM Interface or Subinterface	pvc <vpi/vci> ubr <kbps> (Business Quality Video Streams Support 128kbps - 1Mbps)
Configure ATM Encapsulation Method & Protocol	encapsulation aal5mux ppp virtual-template <n> (n=virtual-template number)
Configure Virtual-template	interface virtual-template <number> (1-25 virtual templates total for all interfaces)
	ip unnumbered ethernet 0/0/0 (or FastEthernet on the NRP)
	peer default ip address pool <word> (the word is optional, example PPPoA)
Configure Global IP Local Pool	ip local pool PPPoA n.n.n.n n.n.n.n
Authentication Method	ppp authentication pap callin (CHAP can also be used)

Enable Multicast PIM on Interface	ip pim sparse-dense-mode Note: ensure you have turned on IP Multicast Routing.
AAA NRP LOCAL PPP Authentication	aaa new-model aaa authentication ppp default local aaa authorization network default local username <name> password <password>
Configure NRP for AAA RADIUS Server Authentication if using Server for Authentication	Radius-server host n.n.n.n auth-port 1645 acct-port 1646 radius-server key <key password>(key secret)
Enable PPP Authentication IPCP Negotiation	Set ppp wan0-0 login <username> Set ppp wan0-0 password<password> Set ppp wan0-0 ipcp 0.0.0.0 Set ppp wan0-0 ipcp subnet 0.0.0.0
Enable DHCP Server	Set dhcp server enabled
Enable NAT (if using private IP)	Set nat enabled

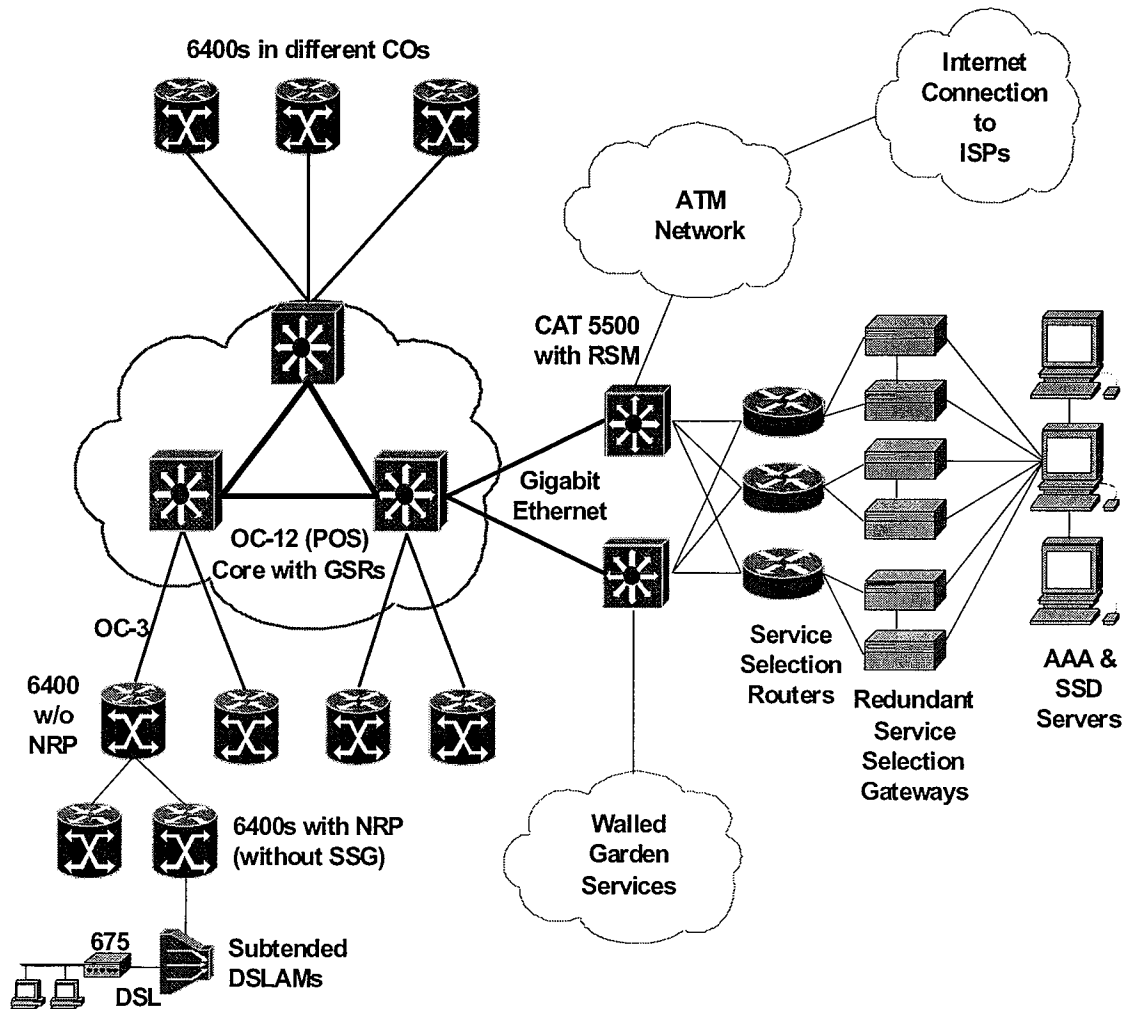
19 Appendix I: NRP Configuration Matrix for RBE

Configure GLOBAL and VPDN-Group COMMANDS (Enable Multicast routing)	ip multicast-routing
Configure ATM Interface and/or Sub-Interface	Atm 0/0/0.n (subinterface) point-to-point ip unnumbered e0/0/0 or fa0/0/0 (or ip addr <address> <mask>) pim sparse-dense-mode
Configure PVCs on ATM Interface or Subinterface	pvc <vpi/vci> ubr <bit rate> Note: 1 PVC per sub-interface
Configure ATM Encapsulation Method & Protocol	Encapsulation aal5snap atm route-bridge ip
Configure Virtual-template	
Configure Global IP Local Pool	
Authentication Method	

Enable Multicast PIM on Interface	ip pim sparse-dense-mode Note: ensure you have turned on IP Multicast Routing.
AAA NRP LOCAL PPP Authentication	
Configure NRP for AAA RADIUS Server Authentication If using Server for Authentication	
Enable PPP Authentication/PCP Negotiation	
Enable DHCP Server	
Enable NAT (if using private IP)	

20 Appendix J: Changing Existing Architectures for Multicast Video Service

This section describes enhancements to some existing network architecture of customers that will enable them to provide multicast video service. The figures show simplified architectures rather than the exact replica. Currently this section describes the enhancements required for one Cisco customer³¹.



The figure above closely follows the architecture deployed by a CLEC who provides DSL access to subscribers. The CLEC (or their subsidiary) also acts as a service provider within their walled garden, and provides several value added services to their customers.

The subscriber premises have Cisco 675 CPE configured for PPP over ATM. The access network uses Cisco 6100 DSLAMs (subtended) and Cisco 6400s. Some of the 6400s are used as switches (without NRP), while the 6400s connected to DSLAMs have NRPs (without the SSG functionality).

³¹ more will be added in future

OC3 uplinks from the DSLAMs provide subscriber aggregation to the 6400s , which in turn, are linked via OC3 to Gigabit Switch Routers. OC12 SONET provides the transport for the GSR backbone, while gigabit ethernet uplinks to Catalyst 5500s give subscribers access to AAA, SSD, Walled Garden Services, and Internet.

For easier management, the AAA servers, SSDs, and SSGs are located in a central site. Note that currently external SSGs (6510s) are used, rather than integrated NRP-SSG. The PPPoA architecture uses VPI/VCI based authentication and IPCP subnet feature for minimizing CPE provisioning while allowing per CPE PPP authentication.

This architecture is similar to the architecture suggested for multicast video service deployment, except that the SSG functionality is provided in this architecture by external 6510 boxes. While this can support multicast video services that do not require authentication, supporting authentication requires the 6400 NRP (that terminates the PPPoA session) to have the SSG functionality. Supporting multicast video service would therefore require removing the external 6510s, and turning on the SSG functionality in the 6400³²s connected to DSLAMs. Other 6400s that work as switches and need not be affected.

Attachments

None

References

1. ATM Half-Bridging Product Requirement Document (ENG -33338)
http://wwwin-eng.cisco.com/Eng/IOS/ATM_Half_Bridging_PRD.doc@latest
2. ATM Half-Bridging Software Unit Functional Specification (ENG-33344)
http://wwwin-eng.cisco.com/Eng/IOS/ATM_Half_Bridging_FS.html@latest
3. xDSL Architecture (ENG-33775)
http://wwwin-eng.cisco.com/Eng/NUBU/xDSL/Solutions_Lab/XDSL_Architecture.doc
4. TopCat/SSG White Paper (ENG-38812)
http://wwwin-eng.cisco.com/Eng/NSMBU/DSL_Provisioning/White_Papers/TopCatSSG-WhitePaper.doc@latest
5. <<Customer name suppressed>> xDSL New Services Architecture (ENG-35100)
6. Tuzigoot Project Plan (ENG-42117)
http://wwwin-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/PP_Tuzigoot.doc
7. VPI-VCI Authenticated Routing AAA Feature Functional Spec" (ENG-33842)
http://wwwin-eng.cisco.com/Eng/NUBU/xDSL/Solutions_Lab/VPI-VCI_Authenticated_Routing_AAA_Feature_FS.doc
8. 6400 RADIUS VC Logging (ENG- 33846)
http://wwwin-eng.cisco.com/Eng/NUBU/xDSL/6400/6400_RADIUS_VC_Logging.txt

³² Supporting multicast authentication with external SSG (6510) would require some kind of IGMP proxy facility on 6400 IOS.

9. Tuzigoot Project Requirement Document (ENG-42271)
http://wwwwin-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/PRD_Tuzigoot.doc
10. Tuzigoot System Functional Specifications (ENG-41377)
http://wwwwin-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/FS_SystemFsTuzigoot.doc
11. Tuzigoot Multicast Authentication Functional Specs ENG-44252
http://wwwwin-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/FS_TuzigootMultiAuthFs.doc
12. For further information on PIM, please refer to
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_c/1cprt1/1cmulti.htm#11494
13. Informational RFC 2516, A Method for Transmitting PPP Over Ethernet (PPPoE), February 1999
14. RFC 1483, Multiprotocol Encapsulation over ATM Adaptation Layer 5, July 1993
15. RFC 2364, PPP over AAL5, July 1998
16. For modification of Cisco Access Registrar (AAA server) to provide VPI/VCI based authentication, please see "Stray Cat System Functional Specification (ENG- 33757)" at:
http://wwwwin-eng.cisco.com/Eng/NSMBU/StrayCat/Sys_Specs/straycatfunc.doc
17. Implementing ADSL, David Ginsburg, Addison-Wesley ISBN 0-201-65760-0

21 End of Document

No information follows this paragraph.

EXHIBIT B

CISCO CONFIDENTIAL

This report contains the following 1 ideas:

Idea No.	Changes	Title	Inventors	Entered	Updated
<u>76761</u>	Method and system for authenticated access to Internet Protocol (IP) Multicast traffic		Sunil Chandrupatla (chandrup),Kali Mishra (kmishra),Sunil Podar (spodar),Sandeep Saksena (ssaksena),Sampath Sthothra Bhasham (ssthothr)		

CISCO CONFIDENTIAL

Method and system for authenticated access to Internet Protocol (IP) Multicast traffic

CPOL No.: 76761 Seq No.: 2797 Status: Pending Submitted: Modified:

Portfolio Manager

Jason Kipnis (jkipnis) Address: Weil Gotshal & Manges LLP
201 Redwood Shores Parkway
Redwood Shores, CA 94065 Email: jason.kipnis@weil.com
Phone:

Idea Details

The contents of this submission and any additions or modifications thereto constitute Cisco confidential information and may be a privileged communication to or from one or more attorneys and/or supporting personnel for purposes of obtaining or facilitating legal advice and/or legal services.

Inventors:

Sunil Chandrupatla (chandrup)	Phone:	Manager: jchristy	Dept: Network Services
Type: Regular	Division:ITD	Site: SAN JOSE SITE 5	Info: --
Kali Mishra (kmishra)	Phone:	Manager: slamarre	Dept: Commercial System...
Type: Regular	Division:CMO Commercial Ma...	Site: RESEARCH TRIANGLE...	Info: --
Sunil Podar (spodar)	Phone:	Manager: along	Dept: Network Services
Type: Regular	Division:ITD	Site: SAN JOSE SITE 5	Info: --
Sandeep Saksena (ssaksena)	Phone:	Manager: vbabu	Dept: NMTG - Campus Mgmt
Type: Regular	Division:NMTG	Site: SAN JOSE SITE 3	Info: --
Sampath Sthothra Bhasham (ssthothr)	Phone:	Manager: shujinz	Dept: MRBU Software
Type: Regular	Division:RSPTG	Site: SAN JOSE SITE 5	Info: --

Background: IP Multicast (IETF RFC 1112) is an Internet standard and a suit of technology that enables distribution of content such as video or audio digital data in a bandwidth-efficient manner through a network. In this technology, IP multicast packets are distributed through a network addressed to a "Multicast Group Address" rather than a destination IP address as in a traditional point-to-point communication. This Multicast

Group Address is not a specific destination and network equipment such as routers understand it as such. The network path that such packets take as they are routed through a network forms a distribution tree. The edge routers (or access routers) receive this data and discard it unless there is some device that has previously expressed an interest in receiving it. A user device that is interested in receiving such data uses another Internet Protocol called Internet Group Management Protocol (IGMP, RFC 2236) and issues a "join" request expressing a desire to receive data of a particular Multicast Group Address. Once the edge router receives such a request, it then forwards the data to the interested device. The advantage of this protocol is that even when there are multiple users interested in receiving the same data, only a single copy of the data travels through the backbone network.

Current technology and standards are geared towards enterprise or corporate networks and permit anyone to "join" a Multicast Group. This enables any user on the network to receive multicast data as long as that data is available on the network. For service provider networks, this poses a problem - what is needed is a way to control access so that only certain users can access only certain content, i.e., controlled privilege to "join" a multicast group. The technology described herein proposes a solution to this problem by providing a mechanism for authenticated access to multicast data.

The basic idea behind authenticated access to IP multicast streams is to intercept the IGMP JOIN request in the access router, authenticate the access privilege for the user via previously provisioned user access information, and depending on the success of authentication, allow or disallow the JOIN request. The present solution uses the RADIUS protocol (RFC 2138 & 2139) for authentication, authorization and accounting. This protocol is well documented and implemented in the dial access solutions. The unique contribution here is the combination of IGMP and RADIUS. A user's access rights are provisioned in a RADIUS server (a.k.a. Authentication, Authorization and Accounting or AAA server). When the router receives a JOIN request, it constructs a RADIUS access-request and forwards to the designated RADIUS server. The RADIUS server replies with an access-accept or access-reject, whereupon the router proceeds to honor or dishonor the JOIN request depending on the reply from the

RADIUS server and undertakes appropriate action to modify its internal routing tables.

We are not aware of any prior solutions to the problem stated above.

Possible Prior Art:**Summary:** < 1st para copied from above>

The basic idea behind authenticated access to IP multicast streams is to intercept the IGMP JOIN request in the access router, authenticate the access privilege for the user via previously provisioned user access information, and depending on the success of authentication, allow or disallow the JOIN request. The present solution uses the RADIUS protocol (RFC 2138 & 2139) for authentication, authorization and accounting. This protocol is well documented and implemented in the dial access solutions. The unique contribution here is the combination of IGMP and RADIUS. A user's access rights are provisioned in a RADIUS server (a.k.a. Authentication, Authorization and Accounting or AAA server). When the router receives a JOIN request, it constructs a RADIUS access-request and forwards to the designated RADIUS server. The RADIUS server replies with an access-accept or access-reject, whereupon the router proceeds to honor or dishonor the JOIN request depending on the reply from the RADIUS server and undertakes appropriate action to modify its internal routing tables.

Per the IGMP standard, when a user issues a JOIN request, only the IP address of the user device is known to a network device receiving the JOIN request. When this idea is used in conjunction with the Cisco Service Selection Gateway technology and product, it enables the network to also identify easily the user who is requesting the JOIN access. It is therefore possible to authenticate the specific user's access privilege rather than simply by the IP address of the accessing device. This is only one way how a username can be associated with the JOIN request. An LDAP directory server can also be used to associate an IP address with a username.

Restatement:

Advantages: Multicast Authentication enables Service Providers to provide controlled access to value-added services based on Multicast content such as video or audio. This capability also enables a subscription-based business model where service providers can bundle different multicast streams into packages of content to which users can subscribe. This opens up a whole new range of revenue opportunities for service providers. Additionally, this makes it possible to have different users have access to different multicast video or audio content; for instance, a parent can subscribe to content different than a child in the same residence.

This is also a unique aspect of the idea presented here since it differentiates the user access model from that of the traditional Cable services. In Cable, it is the device (Set top box) which is the subscriber to premium content, and therefore, enabling the whole household to have the same access privileges.

First Written Record Date:

First Written Record URLs: Multicast Video over DSL Architecture, 1st draft, Section 2.3.3.1, URL:
http://www.win-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/EIN_TuzigootArch.doc@1

Cisco Use: ---Cisco has implemented this feature in its IOS technology and is available currently in a prototype, trial only version on a private branch off of 12.05(DC) version of IOS that is used for the Cisco 6400-NRP-SSG product. This is slated to be used for trials in a few customer trial installations and is targeted to be merged into the mainstream product code in future versions.

Working Model:

Description of possible uses of the technology by others in the industry:

Industry Use:**Public Use:**

---None

Government Use:

Detecting Use: Any company which claims to enable service providers to provide differentiated access to multicast video or audio content where the service provider can bundle the content into packages of services is likely to have implemented multicast authentication on its access router product.

Standards:

Technologies: • IP > IP Multicast > Multicast > Internet Group Management Protocol (IGMP)

- Security and VPN > Authentication Protocols > AAA

Networking Solutions:

- Large Enterprise > Networking Solutions for Large Enterprise > Video Solutions for Large Enterprise > Video Distribution Network Solution
- Large Enterprise > Networking Solutions for Large Enterprise > Access Solutions for Large Enterprise > Internet Access Solution
- Large Enterprise > Networking Solutions for Large Enterprise > Content Networking Solutions for Large Enterprise

Categorization Notes:

Categories Summary
[NM/S&A]

PDDs: ---**Supporting Documents:**

- Delivering Multicast Video over ADSL (White Paper):
- http://wwwin.cisco.com/cmc/cc/cisco/mkt/servprod/dsl/tech/madsl_wp.htm
- Multicast Video over DSL Architecture, Section 9, pp. 39--:
- http://wwwin-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/EIN_TuzigootArch.doc
- Multicast Authentication Functional Specification:
- http://wwwin-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/FS_TuzigootMultiAuthFs.doc
- Multicast Design Document for Tuzigoot:
- http://wwwin-eng.cisco.com/Eng/NUBU/Systems/Projects/DSL_Solutions/Tuzigoot/Specs/DS_MulticastDgn.doc

Documents:

Type	Document	Size
	No Documents Exist.	

Notes: ---

<http://wwwin-eng.cisco.com/pcb/cpol/patent.cgi>

CISCO CONFIDENTIAL

A printed version of this page is an uncontrolled document.

EXHIBIT C

> -----Original Message-----

> From: Sandy Brown <sabrown@cisco.com> [mailto:sabrown@cisco.com]

> Sent:

> To: evertt <evertt@gte.net>; jon.wells <jon.wells@telops.gte.com>

> Cc: robin.gray <robin.gray@telops.gte.com>; lharper
<lharper@cisco.com>;

> spodar <spodar@cisco.com>; epeck <epeck@cisco.com>

> Subject: Executive Briefing

> Robin:

> Please distribute to all the GTE attendees.

> Thanks

> Sandy

> Agenda
> Video Over ADSL Executive Briefing

> 7:22 PM Sandy Brown meets American Flight 432
> and escorts GTE to restaurant

> 8:00 PM Dinner with ClearBand
> 840 North First
> 840 North First Street
> (408) 271-3366
> (Follow directions from San Jose Airport to
> Cisco,below, and stop at 840 North First
Street
> for the
> restaurant)
>

- >
- >
- >
- > Executive Briefing Center
- > Building 10
- > 300 East Tasman Drive
- > San Jose , CA
- >
- > Service Provider Solutions 1 Demo Room
- >
- > 8:30 - 9:00 AM Continental Breakfast
- >
- > 9:00 - 9:10 AM Welcome & Introductions Sandy Brown
- >
- > 9:10 - 10:15 AM Marketing Overview
- > Lance
- > Harper
- > DSL solutions team activities
- >
- > 10:15 - 11:15 AM Network Architecture Overview
- > DSL
- > Solution Team
- >
- > 11:15 - 11:30 AM Break
- >
- > 11:30 AM - 12:30 PM Demonstrations
- > Lance
- > Harper &
- >
- > Solutions Team
- > - Video over ADSL:
- > - ClearBand video server
- > - CoolCast - video content services offering
- > - IP/TV - Cisco solution for VOD
- > - Webcam lower resolution video solution
- > - SSG Web Selection with Dashboard Server
- > - Multicast Authentication
- >
- > 12:30 - 1:00 PM Lunch
- >
- > 1:00 - 2:00 PM ClearBand
- > Presentation Clearband
- > -
- > 2:00 - 2:30 PM Wrap Up
- >
- >
- > Directions from San Jose Airport to
- > Cisco Service Provider EBC Building 10
- >
- > Depart airport towards Highway 101
- > Turn left onto Guadalupe Pkwy.
- > Bear right towards North First Street

- > Turn left onto North First Street
- > Turn right onto Tasman Drive
- > Cross over Zanker (still on Tasman) and turn right
- > on East
- > Morgridge Way, entering building 10 at 300
- > Tasman
- > Visitor Parking in front
- > Approximate distance from airport is 5 miles
- >
- >
- > GTE Attendees
- > Evertt Williams, Vice President, National Data Marketing Management
- > Jon Wells, Director, Data Development Strategy Matt Daniels, Manager,
- > Data Development Ron Zilolkowski, Director, Business Development
- > Janet Sun, Assistant Vice President, Data Development
- > Mark Dillion, Director, Applications Development
- > Ernie Levinson, Group Manager, Application, Development
- > Arthur Solozar, Manager, Network Architecture
- > Paul Symczak, Director, Strategic Alliances
- > Cisco Systems Empowering the Internet Generation
- >